AZTCOFW

The AZTCO-FW Documentation

www.aztcofirewall.sy

AZTCO

FEB 21, 2021

Contents

ONE

PREFACE

1.1 Acknowledgements

AZTCO-FW was graded a Visionary in the Gartner Magic Quadrant. It provides next-generation firewall protection that's relatively easy to set up and manage. It blocks unknown threats, automatically responds to security incidents by isolating compromised systems, and exposes hidden user, application and threat risks on the network. AZTCO-FW also includes a web application firewall, ransomware protection, phishing prevention, all firewall rules unified on a single screen, and a secure web gateway.

AZTCO-FW is based on AZTCO-FW Opensource Firewall that is customized and developed by Syrian developers' hands.

1.2 Feedback

Welcome to *The AZTCO-FW Documentation*, written by the AZTCO-FW project team

This set of documents covers topics ranging from the installation process and basic configuration to advanced net- working and firewalling

This is designed to be a friendly guide to common networking and security tasks along with a thorough reference for the capabilities of AZTCO-FW. These documents cover the following topics (and more!):

An introduction to AZTCO-FW and its features.

Firewall design and hardware planning.

Installing and upgrading AZTCO-FW.

Using the web-based configuration interface.

Backing up and restoring the firewall configuration

Firewalling fundamentals including defining and troubleshooting rules.

Port forwarding and Network Address Translation (NAT).

General networking and routing configuration.

Virtual LANs (VLANs), Multi-WAN, and Bridging.

Virtual Private Networks using IPsec and OpenVPN.

Traffic shaping using ALTQ or Limiters.

Captive Portal setup.

High Availability using redundant firewalls.

Various network-related services.

Firewall monitoring, logging, traffic analysis, sniffing, packet capturing, and troubleshooting.

IDS/IPS configuration

Proxy configuration

For general feedback related to the AZTCO-FW project, please post to the forum. Links to these resources can be found at https://aztcofirewall.com/feedback/

CHAPTER

TWO

INTRODUCTION

2.1 Common Deployments

AZTCO-FW® software can meet the needs of nearly any type and size of network environment, from a SOHO to datacenter environments. This section outlines the most common deployments.

2.1.1 Perimeter Firewall

The most common deployment of AZTCO-FW software is a perimeter firewall. AZTCO-FW accommodates networks requiring multiple Internet connections, multiple LAN networks, and multiple DMZ networks. BGP (Border Gateway Protocol), connection redundancy, and load balancing capabilities are configurable as well.

See also:

These advanced features are further described in *Routing* and *Multiple WAN Connections*.

2.1.2 LAN or WAN Router

AZTCO-FW software configured as a LAN or WAN router and perimeter firewall is a common deployment in small networks. LAN and WAN routing are separate roles in larger networks.

LAN Router

AZTCO-FW software is a proven solution for connecting multiple internal network segments. This is most commonly deployed with VLANs configured with 802.1Q trunking, described more in *Virtual LANs (VLANs)*. Multiple Ethernet interfaces are also used in some environments. High-volume LAN traffic environments with fewer filtering requirements may need layer 3 switches or ASIC-based routers instead.

WAN Router

AZTCO-FW is a great solution for Internet Service Providers. It offers all the functionality required by most networks at a much lower price point than other commercial offerings.

2.1.3 Special Purpose Appliances

AZTCO-FW can be utilized for less common deployment scenarios as a stand-alone appliance. Examples include: VPN appliance, Sniffer appliance, and DHCP server appliance.

VPN Appliance

AZTCO-FW software installed as a separate Virtual Private Network appliance adds VPN capabilities without disrupting the existing firewall infrastructure, and includes multiple VPN protocols.

DHCP Server Appliance

AZTCO-FW software can be deployed strictly as a Dynamic Host Configuration Protocol server, however, there are limitations of the AZTCO-FW GUI for advanced configuration of the ISC DHCP daemon.

See also:

For more information on configuring the DHCP service on AZTCO-FW, see *DHCP*.

2.2 Interface Naming Terminology

All interfaces on a AZTCO-FW® router/firewall can be assigned any name desired, but they all start with default names: WAN, LAN, and OPT.

2.2.1 WAN

Short for *Wide Area Network*, WAN is the untrusted public network outside of the firewall. In other words, the WAN interface is the firewall's connection to the Internet or other upstream network. In a multi-WAN deployment, WAN is the first or primary Internet connection.

At a minimum, the firewall must have one interface, and that is WAN.

2.2.2 LAN

Short for *Local Area Network*, LAN is commonly the private side of a firewall. It typically utilizes a *private IP address* scheme for local clients. In small deployments, LAN is typically the only internal interface.

2.4.3 OPT

OPT or *Optional* interfaces refer to any additional interfaces other than WAN and LAN. OPT interfaces can be additional LAN segments, WAN connections, DMZ segments, interconnections to other private networks, and so on.

2.4.4 DMZ

Short for the military term *demilitarized zone*, DMZ refers to the buffer between a protected area and a war zone. In networking, it is an area where public servers are reachable from the Internet via the WAN but isolated from the LAN. The DMZ keeps the systems in other segments from being endangered if the network is compromised, while also protecting hosts in the DMZ from other local segments and the Internet in general.

Warning: Some companies misuse the term “DMZ” in their firewall products as a reference to 1:1 NAT on the WAN IP address which exposes a host on the LAN.

2.3 Finding Information and Getting Help

This section offers guidance on finding information in this documentation, on AZTCO-FW® software in general, as well as providing further resources.

2.3.1 Getting Help

A help icon is available on almost every page, ❓, and links to the associated page in documentation.

AZTCO-FW also offer supporting via support form can found on: https://aztcofirewall.com/support/

THREE

CONFIGURATION

3.1 Setup Wizard

The first time a user logs into the AZTCO-FW® software GUI, the firewall presents the Setup Wizard automatically. The first page of the wizard

Click » Next to proceed.

Tip: Using the setup wizard is optional. Click the logo at the top left of the page to exit the wizard at any time.

The next screen of the wizard explains the availability of support from AZTCO-FW. Click » Next again to start the configuration process using the wizard.

3.1.1 General Information Screen

The next screen (Figure *General Information Screen*) configures the name of this firewall, the domain in which it resides, and the DNS servers for the firewall.

Hostname the Hostname is a name that should uniquely identify this firewall. It can be nearly anything, but must start with a letter and it may contain only letters, numbers, or a hyphen.

Domain Enter a Domain, e.g., example.com. If this network does not have a domain, use <something>. home. arpa, where <something> is another identifier: a company name, last name, nickname, etc. For example, company.home.arpa the hostname and domain name are combined to make up the fully qualified domain name of this firewall.

Primary/Secondary DNS Server The IP address of the Primary DNS Server and Secondary DNS Server, if known.

These DNS servers may be left blank if the DNS Resolver will remain active using its default settings. The default configuration has the DNS Resolver active in resolver mode (not forwarding mode), when set this way, the DNS Resolver does not need forwarding DNS servers as it will communicate directly with Root DNS servers and other authoritative DNS servers. To force the firewall to use these configured DNS servers, enable forwarding mode in the DNS Resolver or use the DNS Forwarder.

If this firewall has a dynamic WAN type such as DHCP, PPTP or PPPoE these may be automatically assigned by the ISP and can be left blank.

Note: The firewall can have more than two DNS servers, add more under System > General Setup after completing the wizard.

Override DNS When checked, a dynamic WAN ISP can supply DNS servers which override those set manually. To force the use of only the DNS servers configured manually, uncheck this option.

See also:

For more information on configuring the DNS Resolver, see *DNS Resolver*

Click » Next to continue.

3.1.2 NTP and Time Zone Configuration

The next screen (Figure *NTP and Time Zone Setup Screen*) has time-related options.

Time server hostname A Network Time Protocol (NTP) server hostname or IP address. Unless a specific NTP server is required, such as one on LAN, the best practice is to leave the Time server hostname at the default time.aztcofirewall.sy. This value will pick a set of random servers from a pool of known-good NTP hosts.

To utilize multiple time server pools or individual servers, add them in the same box, separating each server by a space. For example, to use three NTP servers from the pool, enter: 0. time.aztcofirewall.sy 1. time.aztcofirewall.sy 2. time.aztcofirewall.sy

This numbering is specific to how.pool.ntp.org operates and ensures each address is drawn from a unique pool of NTP servers so the same server does not get used twice.

Time zone Choose a geographically named zone which best matches location of this firewall, or any other desired zone.

Click » Next to continue.

3.1.3 WAN Configuration

The next page of the wizard configures the WAN interface of the firewall. This is the external network facing the ISP or upstream router, so the wizard offers configuration choices to support several common ISP connection types.

WAN Type the Selected Type (Figure *WAN Configuration*) must match the type of WAN required by the ISP, or whatever the previous firewall or router was configured to use. Possible choices are *Static*, *DHCP*, *PPPoE*, and *PPTP*. The default choice is *DHCP* due to the fact that it is the most common, and for the majority of cases this setting allows a firewall to “Just Work” without additional configuration. If the WAN type is not known, or specific settings for the WAN are not known, this information must be obtained from the ISP. If the required WAN type is not available in the wizard,

or to read more information about the different WAN types, see *Interface Types and Configuration*.

Note: If the WAN interface is wireless, additional options will be presented by the wizard which are not covered during this walkthrough of the standard Setup Wizard. Refer to *Wireless*, which has a section on Wireless WAN for additional information. If any of the options are unclear, skip the WAN setup for now, and then perform the wireless configuration afterward.

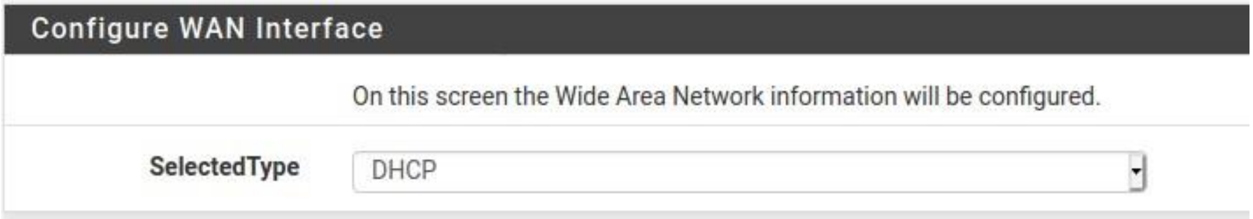

Configure WAN Interface

On this screen the Wide Area Network information will be configured.

SelectedType: DHCP

Fig. 4: WAN Configuration

MAC Address This field, shown in Figure *General WAN Configuration*, changes the MAC address used on the WAN network interface. This is also known as "spoofing" the MAC address.

Note: The problems alleviated by spoofing a MAC address are typically temporary and easily worked around. The best course of action is to maintain the original hardware MAC address, resorting to spoofing only when absolutely necessary.

Changing the MAC address can be useful when replacing an existing piece of network equipment. Certain ISPs, primarily Cable providers, will not work properly if a new MAC address is encountered. Some Internet providers require power cycling the modem, others require registering the new address over the phone. Additionally, if this WAN connection is on a network segment with other systems that locate it via ARP, changing the MAC to match and older piece of equipment may also help ease the transition, rather than having to clear ARP caches or update static ARP entries.

Warning: If this firewall will ever be used as part of a *High Availability Cluster*, do not spoof the MAC address.

Maximum Transmission Unit (MTU) The MTU field, shown in Figure *General WAN Configuration*, can typically be left blank, but can be changed when necessary. Some situations may call for a lower MTU to ensure packets are sized appropriately for an Internet connection. In most cases, the default assumed values for the WAN connection type will work properly.

Maximum Segment Size (MSS) MSS, shown in Figure *General WAN Configuration* can typically be left blank, but can be changed when necessary. This field enables MSS clamping, which ensures TCP packet sizes remain adequately small for a particular Internet connection.

Static IP Configuration If the "Static" choice for the WAN type is selected, the IP address, Subnet Mask, and Upstream Gateway must all be filled in (Figure *Static IP Settings*). This information must be obtained from the ISP or whoever controls the network on the WAN side of this firewall. The IP Address and Upstream Gateway must both reside in the same Subnet.

DHCP Hostname This field (Figure *DHCP Hostname Setting*) is only required by a few ISPs. This value is sent along with the DHCP request to obtain a WAN IP address. If the value for this field is unknown, try leaving it blank unless directed otherwise by the ISP.

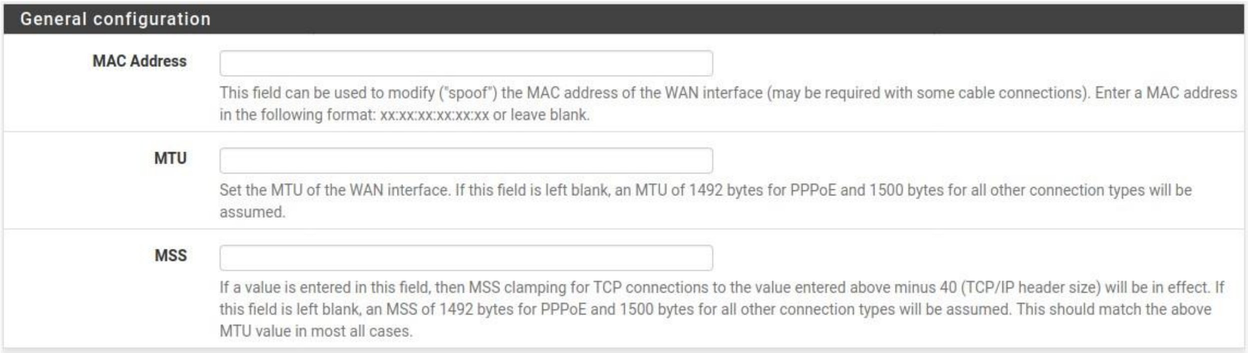

General configuration

MAC Address

This field can be used to modify ("spoof") the MAC address of the WAN interface (may be required with some cable connections). Enter a MAC address in the following format: xx:xx:xx:xx:xx:xx or leave blank.

MTU

Set the MTU of the WAN interface. If this field is left blank, an MTU of 1492 bytes for PPPoE and 1500 bytes for all other connection types will be assumed.

MSS

If a value is entered in this field, then MSS clamping for TCP connections to the value entered above minus 40 (TCP/IP header size) will be in effect. If this field is left blank, an MSS of 1492 bytes for PPPoE and 1500 bytes for all other connection types will be assumed. This should match the above MTU value in most all cases.

Fig. 5: General WAN Configuration

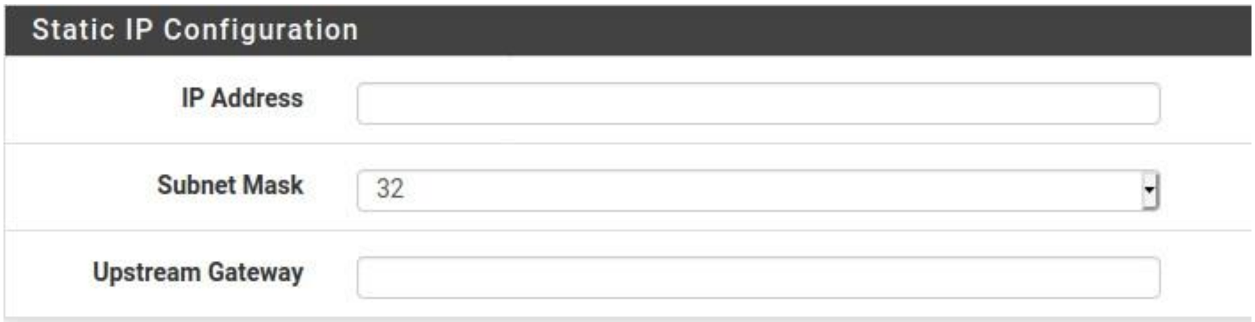

Static IP Configuration

IP Address

Subnet Mask: 32

Upstream Gateway

Fig. 6: Static IP Settings

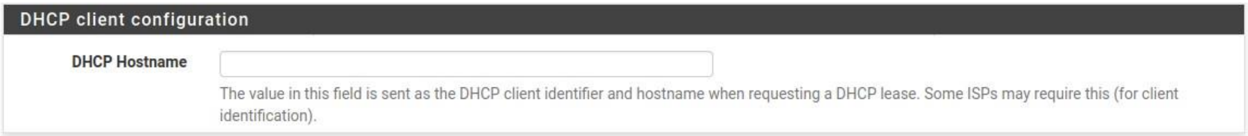

DHCP client configuration

DHCP Hostname

The value in this field is sent as the DHCP client identifier and hostname when requesting a DHCP lease. Some ISPs may require this (for client identification).

Fig. 7: DHCP Hostname Setting PPPoE Configuration When using the PPPoE (Point-to-Point Protocol over Ethernet) WAN type (Figure *PPPoE Configuration*), The PPPoE Username and PPPoE Password fields are required, at a minimum. The values for these fields are determined by the ISP.

PPPoE Username The login name for PPPoE authentication. The format is controlled by the ISP, but commonly uses an e-mail address style such as myname@example.com.

PPPoE Password The password to login to the account specified by the username above. The password is masked by default. To view the entered password, check Reveal password characters.

PPPoE Service Name the PPPoE Service name may be required by an ISP, but is typically left blank. When in doubt, leave it blank or contact the ISP and ask if it is necessary.

PPPoE Dial on Demand This option leaves the connection down/offline until data is requested that would need the connection to the Internet. PPPoE logins happen quite fast, so in most cases the delay while the connection is setup would be negligible. If public services are hosted behind this firewall, do not check this option as an online connection must be maintained as much as possible in that case. Also note that this choice will not drop an existing connection.

PPPoE Idle Timeout Specifies how much time the PPPoE connection remain up without transmitting data before disconnecting. This is only useful when coupled with Dial on demand, and is typically left blank (disabled).

Note: This option also requires the deactivation of gateway monitoring, otherwise the connection will never be idle.

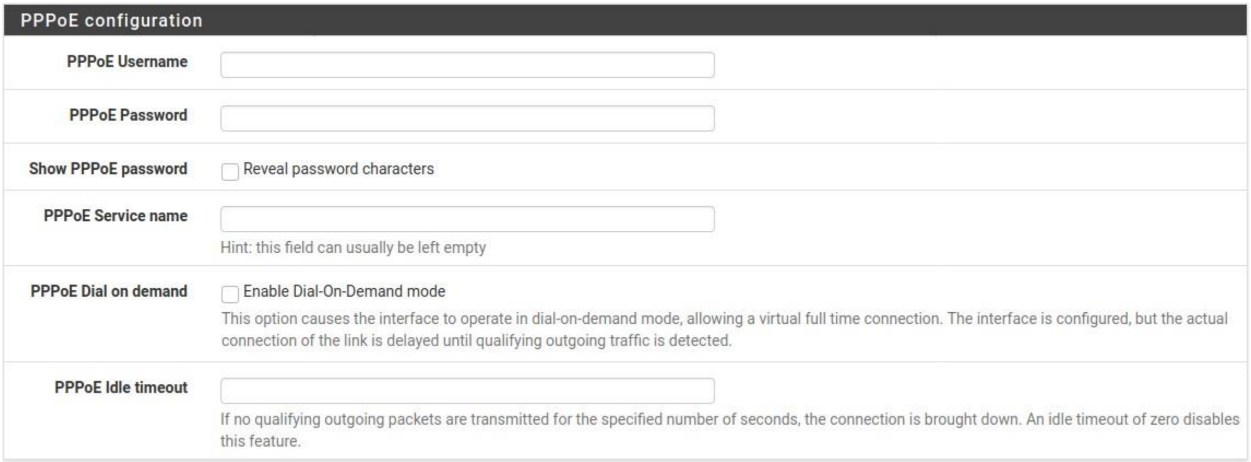

PPPoE configuration

PPPoE Username

PPPoE Password

Show PPPoE password □ Reveal password characters

PPPoE Service name

Hint: this field can usually be left empty

PPPoE Dial on demand □ Enable Dial-On-Demand mode

This option causes the interface to operate in dial-on-demand mode, allowing a virtual full time connection. The interface is configured, but the actual connection of the link is delayed until qualifying outgoing traffic is detected.

PPPoE Idle timeout

If no qualifying outgoing packets are transmitted for the specified number of seconds, the connection is brought down. An idle timeout of zero disables this feature.

Fig. 8: PPPoE Configuration

PPTP Configuration the PPTP (Point-to-Point Tunneling Protocol) WAN type (Figure *PPTP WAN Configuration*) is for ISPs that require a PPTP login, not for connecting to a remote PPTP VPN. These settings, much like the PPPoE settings, will be provided by the ISP. A few additional options are required:

Local IP Address The local (usually private) address used by this firewall to establish the PPTP connection.

CIDR Subnet Mask The subnet mask for the local address.

Remote IP Address the PPTP server address, which is usually inside the same subnet as the Local IP address.

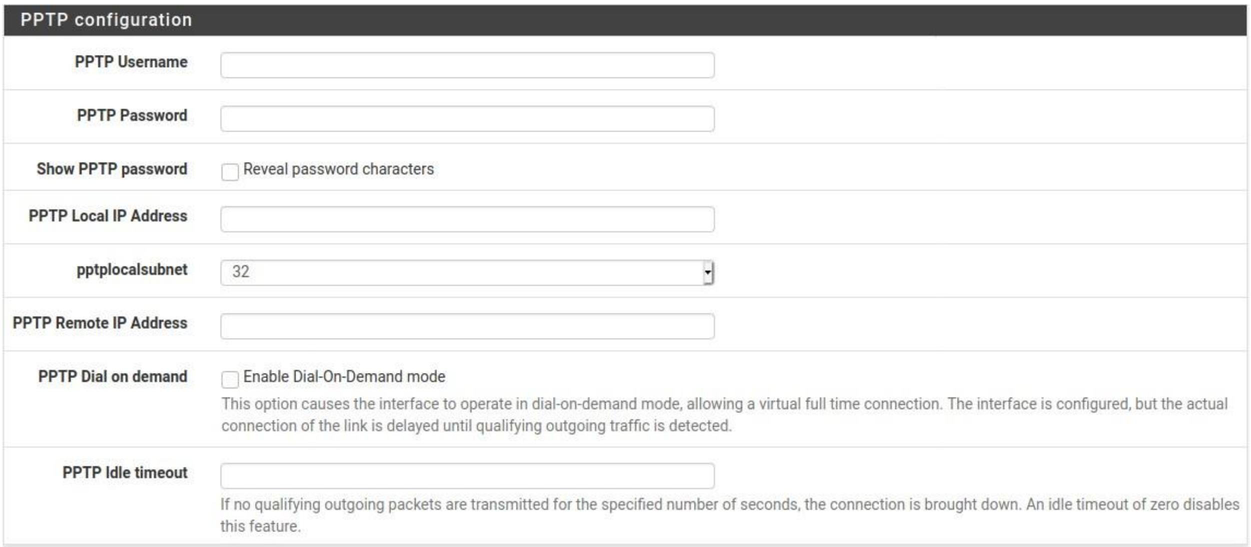

PPTP configuration

PPTP Username	
PPTP Password	
Show PPTP password	□ Reveal password characters
PPTP Local IP Address	
pptplocalsubnet	32
PPTP Remote IP Address	
PPTP Dial on demand	□ Enable Dial-On-Demand mode This option causes the interface to operate in dial-on-demand mode, allowing a virtual full time connection. The interface is configured, but the actual connection of the link is delayed until qualifying outgoing traffic is detected.
PPTP Idle timeout	If no qualifying outgoing packets are transmitted for the specified number of seconds, the connection is brought down. An idle timeout of zero disables this feature.

Fig. 9: PPTP WAN Configuration

These last two options, seen in Figure *Built-in Ingress Filtering Options*, are useful for preventing invalid traffic from entering the network protected by this firewall, also known as "Ingress Filtering".

Block RFC 1918 Private Networks Blocks connections sourced from registered private networks such as 192.168.x.x and 10.x.x.x attempting to enter the WAN interface. A full list of these networks is in *Private IP Addresses*.

Block Bogon Networks When active, the firewall blocks traffic from entering if it is sourced from reserved or unassigned IP space that should not be in use. The list of bogon networks is updated periodically in the background, and requires no manual maintenance. Bogon networks are further explained in *Block Bogon Networks*.

Click » Next to continue once the WAN settings have been filled in.

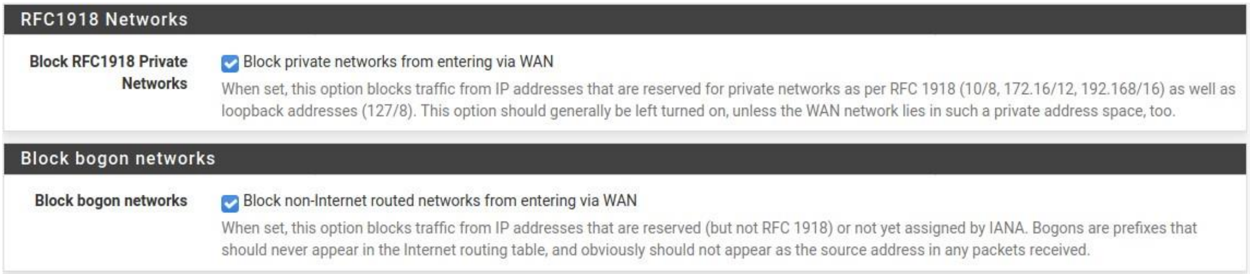

RFC1918 Networks

Block RFC1918 Private Networks	☑ Block private networks from entering via WAN When set, this option blocks traffic from IP addresses that are reserved for private networks as per RFC 1918 (10/8, 172.16/12, 192.168/16) as well as loopback addresses (127/8). This option should generally be left turned on, unless the WAN network lies in such a private address space, too.

Block bogon networks

Block bogon networks	☑ Block non-Internet routed networks from entering via WAN When set, this option blocks traffic from IP addresses that are reserved (but not RFC 1918) or not yet assigned by IANA. Bogons are prefixes that should never appear in the Internet routing table, and obviously should not appear as the source address in any packets received.

Fig. 10: Built-in Ingress Filtering Options

3.1.4 LAN Interface Configuration

This page of the wizard configures the LAN IP Address and Subnet Mask (Figure *LAN Configuration*).

If this firewall will not connect to any other network via VPN, the default 192.168.1.0/24 network may be acceptable. If this network must be connected to another network, including via VPN from remote locations, choose a private IP address range much more obscure than the common default of 192.168.1.0/24. IP space within the 172.16.0.0/12

RFC 1918 private address block is generally the least frequently used, so choose something between 172.16.x.x and 172.31.x.x to help avoid VPN connectivity difficulties.

If the LAN is 192.168.1.x and a remote client is at a wireless hotspot using 192.168.1.x (very common), the client will not be able to communicate across the VPN. In that case, 192.168.1.x is the local network for the client at the hotspot, not the remote network over the VPN.

If the LAN IP Address must be changed, enter it here along with a new Subnet Mask. If these settings are changed, the IP address of the computer used to complete the wizard must also be changed if it is connected through the LAN. Release/renew its DHCP lease, or perform a "Repair" or "Diagnose" on the network interface when finished with the setup wizard.

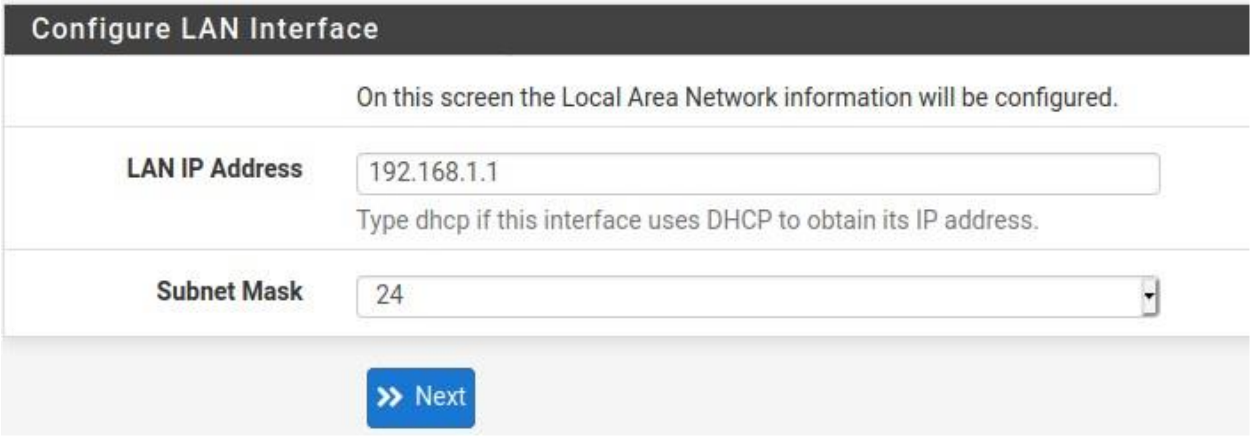

Configure LAN Interface

On this screen the Local Area Network information will be configured.

LAN IP Address: 192.168.1.1

Type dhcp if this interface uses DHCP to obtain its IP address.

Subnet Mask: 24

Fig. 11: LAN Configuration

Click » Next to continue.

3.1.5 Set admin password

Next, change the administrative password for the GUI as shown in Figure *Change Administrative Password*. The best practice is to use a strong and secure password, but no restrictions are automatically enforced. Enter the password in the Admin Password and confirmation box to be sure that has been entered correctly.

Click » Next to continue.

Warning: Do not leave the password set to the default aztco. If access to the firewall administration via GUI or SSH is exposed to the Internet, intentionally or accidentally, the firewall could easily be compromised if it still uses the default password.

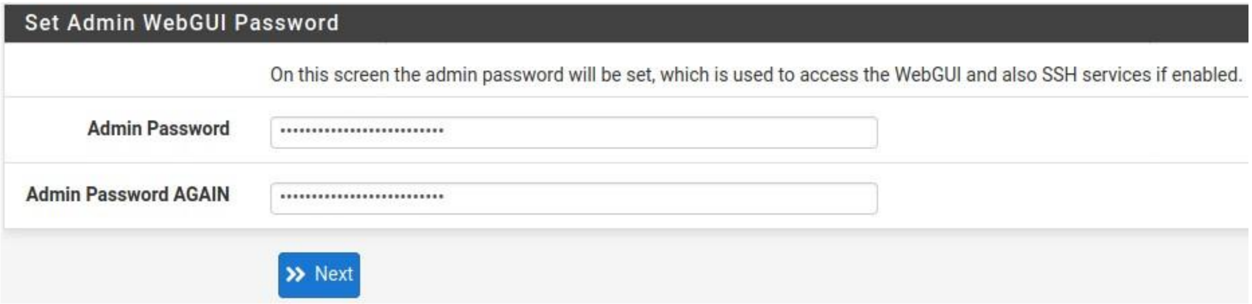

Set Admin WebGUI Password

On this screen the admin password will be set, which is used to access the WebGUI and also SSH services if enabled.

Admin Password ••••••••••••••••••••••••••••

Admin Password AGAIN ••••••••••••••••••••••••••••

Fig. 12: Change Administrative Password

3.1.6 Completing the Setup Wizard

That completes the setup wizard configuration. Click Reload (Figure *Reload the GUI*) and the GUI will apply the settings from the wizard and reload services changed by the wizard.

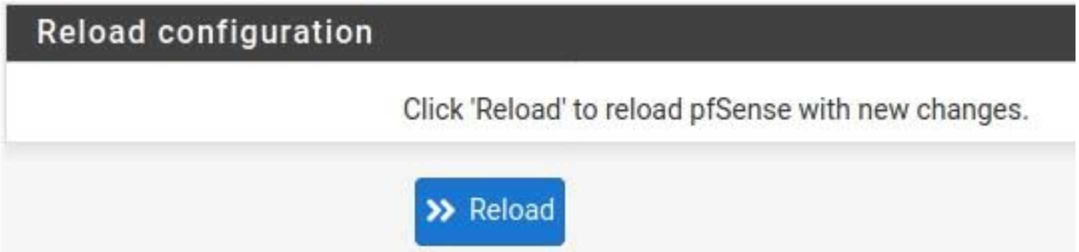

Reload configuration

Click 'Reload' to reload pfSense with new changes.

Fig. 13: Reload the GUI

Tip: If the LAN IP address was changed in the wizard and the wizard was run from the LAN, adjust the client computer's IP address accordingly after clicking Reload.

When prompted to login again, enter the new password. The username remains admin.

After reloading, the final screen of the wizard includes convenient links to check for updates, get support, and other resources. Click Finish to complete and exit the wizard.

At this point the firewall will have basic connectivity to the Internet via the WAN and clients on the LAN side will be able to reach Internet sites through this firewall.

If at any time this initial configuration must be repeated, revisit the wizard at System > Setup Wizard from within the GUI.

3.2 Interface Configuration

Basic aspects of interface configuration within AZTCO-FW® software can be performed at the console and in the setup wizard to start, but changes may also be made after the initial setup by visiting pages under the Interfaces menu. A few basics are covered here, the details can be found in *Interface Types and Configuration*.

8.2. Interface Configuration

3.2.1 Assign interfaces

Interfaces added after the initial setup may be assigned roles by visiting Interfaces > Assignments. There are numerous tabs on that page used for assigning and creating different types of interfaces. The two most commonly used tabs are Interface assignments and VLANs.

See also:

VLAN configuration is covered in *Virtual LANs (VLANs)*.

The Interface assignments tab shows a list of all currently assigned interfaces: WAN, LAN, and any OPTx entries configured on the firewall. Next to each interface is a drop-down list of all network interfaces/ports found on the system. This list includes hardware interfaces as well as VLAN interfaces and other virtual interface types. The MAC address, VLAN tag, or other identifying information is printed alongside the interface name to aid in identification.

The other tabs, much like the VLAN tab, are there to create additional interfaces which can then be assigned. All of these interface types are covered in *Interface Types and Configuration*.

To change an existing interface assignment to another network port:

- Navigate to Interfaces > Assignments
- Locate the interface to change in the list
- Select the new network port from the drop-down list on the row for that interface
- Click Save

To add a new interface from the list of unused network ports:

- Navigate to Interfaces > Assignments
- Select the port to use from the drop-down list labeled Available Network Ports
- Click ➕ Add

This action will add another line with a new OPT interface numbered higher than any existing OPT interface, or if this is the first additional interface, *OPT1*.

3.2.2 Interface Configuration Basics

Interfaces are configured by choosing their entry from under the Interfaces menu. For example, to configure the WAN interface, choose Interfaces > WAN.

Every interface is configured in the same manner and any interface can be configured as any interface type (Static, DHCP, PPPoE, etc.). Additionally, the blocking of private networks and bogon networks may be performed on any interface. Every interface can be renamed, including WAN and LAN, to a custom name. Furthermore, every interface can be enabled and disabled as desired, so long as a minimum of one interface remains enabled.

See also:

For detailed interface configuration information, see *Interface Types and Configuration*

The IPv4 Configuration Type can be changed between *Static IPv4*, *DHCP*, *PPPoE*, *PPP*, *PPTP*, *L2TP*, or *None* to leave the interface without an IPv4 address. When *Static IPv4* is used, an IPv4 Address, subnet mask, and IPv4 Upstream Gateway may be set. If one of the other options is chosen, then type-specific fields appear to configure each type.

The IPv6 Configuration Type can be set to *Static IPv6*, *DHCP6*, *SLAAC*, *6rd Tunnel*, *6to4 Tunnel*, *Track Interface*, or *None* to leave IPv6 unconfigured on the interface. When Static IPv6 is selected, set an IPv6 address, prefix length, and IPv6 Upstream Gateway.

8.2. Interface Configuration

If this a wireless interface, the page will contain many additional options to configure the wireless portion of the interface. Consult *Wireless* for details.

Note: Selecting a Gateway from the drop-down list, or adding a new gateway and selecting it, will direct the firewall to treat this interface as a WAN type interface for NAT and related functions. This is not desirable for internal-facing interfaces such as LAN or a DMZ. Gateways may still be utilized on those interfaces for static routes and other purposes *without* selecting a Gateway here on the interfaces page.

3.3 Managing Lists in the GUI

The AZTCO-FW® software GUI has a common set of icons which are used for managing lists and collections of objects throughout the firewall. Not every icon is used in every page, but their meanings are consistent based on the context in which they are seen. Examples of such lists include firewall rules, NAT rules, IPsec, OpenVPN, and certificates.

➕ Add a new item to a list

↥ Add an item to the beginning of a list

↧ Add an item to the end of a list

✎ Edit an existing item

⧉ Copy an item (create a new item based on the selected item)

⊘ Disable an active item

☑ Enable a disabled item

🗑 Delete an item

⚓

Used for moving entries after selecting one or more items. Click to move the selected items above this row. Shift-click to move the selected items below this row.

Sections may have their own icons specific to each area. Consult the appropriate sections of this documentation for specifics about icons found in other parts of the firewall.

3.4 Quickly Navigate the GUI with Shortcuts

Many areas of the GUI have shortcut icons present in the area known as the "Breadcrumb Bar", as seen in Figure *Shortcuts Example*. These shortcut icons reduce the amount of hunting required to locate related pages, allowing a firewall administrator to navigate quickly between the status page of a service, its logs, and configuration. The shortcuts for a given topic are present on each page related to that topic.

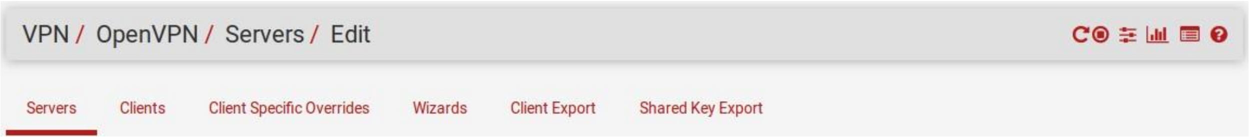

Fig. 14: Shortcuts Example

Note: Shortcut icons only appear when their respective actions are possible and the target pages exist. Not every section has every icon.

The shortcut icons have the following effects when they appear in the GUI:

Start Service If the service is stopped, this icon starts the service.

Restart Service If the service is running, this icon restarts the service.

Note: Some services will stop and start, others reload the configuration. Check the documentation of each service for details.

Stop Service If the service is running, this icon stops the service.

Related Settings This icon navigates to the settings page for this section.

Status Page Link This icon navigates to the the status page for this section.

Log Page Link This icon navigates to the the logs page for this section.

Help Link This icon navigates to a related help topic for this page.

The *Service Status* page (Status > Services) also has shortcut controls for pages related to each service, as shown in Figure *Shortcuts on Service Status*. The icons have the same meaning as in the above section.

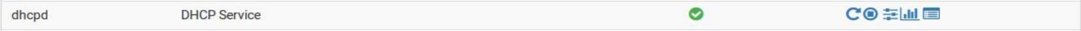

dhcpd DHCP Service

Fig. 15: Shortcuts on Service Status

3.5 General Configuration Options

System > General Setup contains options which set basic configuration items for AZTCO-FW® software. A few of these options are also found in the *Setup Wizard*.

Hostname the Hostname is the short name for this firewall, such as firewall1, hq-fw, or site1. The name must start with a letter and it may contain only letters, numbers, or a hyphen.

Domain the Domain name for this firewall, e.g., example.com. If this network does not have a domain, use <something>. localdomain, where <something> is another identifier: a company name, last name, nickname, etc. For example, company.localdomain

The Hostname and Domain name are combined to make up the Fully Qualified Domain Name (FQDN) of this firewall. For example, if the Hostname is fw1 and the Domain is example.com, then the FQDN is fw1.example. com.

3.5.1 DNS Server Settings

Options in this section control how the firewall resolves hostnames using DNS.

Note: The DNS Resolver is active by default and uses resolver mode (*DNS Resolver*). When set this way the DNS Resolver does not need forwarding DNS servers as it will communicate directly with root DNS servers and other authoritative DNS servers.

To use the servers in this list, switch the DNS resolver to forwarding mode. The DNS Forwarder (*DNS Forwarder*) only supports forwarding mode and will always use the servers from this list.

DNS Servers

This page supports multiple DNS servers managed as a list. To add more DNS servers, click ✚ Add DNS Server.

To remove an entry from the list, click 🗑 Delete.

The DNS server list may be left blank if the DNS Resolver will remain active using its default settings. If this firewall has a dynamic WAN type such as DHCP, PPTP or PPPoE these may be automatically assigned by the ISP and can also be left blank.

Each DNS server entry has the following properties:

DNS Server Address the IP address of the DNS Server.

DNS Server Hostname the FQDN of the DNS server, used to validate DNS server certificates when using DNS over TLS (*DNS Resolver*).

DNS Server Gateway The gateway through which the firewall will reach this DNS server.

This is useful in a Multi-WAN scenario where, ideally, the firewall will have at least one DNS server configured per WAN. More information on DNS for Multi- WAN can be found in *DNS Servers and Static Routes*.

DNS Resolution Behavior

These options fine tune the way the firewall utilizes DNS servers.

DNS Server Override When checked, a dynamic WAN ISP can supply DNS servers which override those set manually. To force the use of only the DNS servers on this page, uncheck this option. This does not apply to the DNS Resolver when acting in resolver mode.

Disable DNS Forwarder By default, the firewall will consult the DNS Resolver or DNS Forwarder running on this firewall to resolve hostnames for itself. It does this by listing localhost (127.0.0.1) as its first DNS server internally. Activating this option disables this behavior, forcing the firewall to use the DNS servers configured above instead of itself.

3.5.2 Localization

Options in this section control the firewall’s clock display and language.

Time Zone The time zone used by the firewall for its clock. Choose a geographically named zone which best matches location of this firewall, or a common zone such as UTC. The firewall clock, log entries, and other areas of the firewall base their time on this zone. Changing the zone requires a reboot to fully activate the new zone in all areas of the firewall.

Time Servers *Network Time Protocol (NTP)* server hostnames or IP addresses. Unless a specific NTP server is required, such as one on LAN, the best practice is to leave the Time Servers value at the default 0.AZTCO-FW.pool.ntp.org. This value will pick random servers from a pool of known-good NTP hosts.

To utilize multiple time servers or pools, add them in the same box, separating each entry by a space. For example, to use three NTP servers from the pool, enter:

```
0.AZTCO-FW.pool.ntp.org 1.AZTCO-FW.pool.ntp.org 2.AZTCO-FW.pool.ntp.org
```

This numbering is specific to how.pool.ntp.org operates and ensures each address is drawn from a unique pool of NTP servers so the same server does not get used twice.

Language The language used by the GUI. The GUI has been translated into multiple languages in addition to the default *English* language.

3.5.3 webConfigurator

Options in this section control various aspects of GUI behavior.

Theme the Theme controls the look and feel of the GUI. Several themes are included in the base system, and they only make cosmetic not functional changes to the WebGUI.

Top Navigation This option controls the behavior of the menu bar at the top of each page. There are two possible choices:

Scrolls with page the default behavior. When the page scrolls, the navigation remains at the top of the *page*, so it is no longer visible as it scrolls off the top of the window. This is the best option for most situations.

Fixed When selected, the navigation remains fixed at the top of the *window*, always visible and available for use. This behavior can be convenient, but on smaller screens such as tablets and mobile devices, long menus can be cut off, leaving options at the bottom unreachable.

Hostname in Menu When set, the GUI includes the firewall Hostname or Fully Qualified Domain Name in the menu bar for reference. This can aid when maintaining multiple firewalls, making it easier to distinguish them without looking at the browser title or tab text.

Dashboard Columns The dashboard is limited to 2 columns by default. On wider displays, additional columns can make better use of horizontal screen space. The maximum number of columns is 4.

Interfaces Sort When unset (default), the GUI presents interfaces in their natural order from the configuration. This is critical for functions such as High Availability which require specific interface ordering. When this option is set, the GUI sorts the interface list alphabetically.

Associated Panels Show/Hide Some GUI pages contain collapsible panels with settings or functions. These panels take up extra screen space, so they are hidden by default. For firewall administrators who use the panels frequently, this can be slow and inefficient. The options in this group make the GUI show these panels by default instead of hiding them.

Available Widgets Controls the Available Widgets panel on the Dashboard.

Log Filter Controls the log filtering (▼) panel used for searching log entries under Status > System Logs.

Manage Log Controls the per-log settings in the Manage Log (🔧) panel available for each log under Status > System Logs.

Monitoring Settings Controls the options panel used to change the graphs at Status > Monitoring.

Require State Filter When set, the state table contents at Diagnostics > States are suppressed by the GUI unless a filter string is present. This helps the GUI handle large state tables which otherwise may fail to load.

Left Column Labels When checked, the option labels in the left column are set to toggle options when clicked. This can be convenient if the firewall administrator is used to the behavior, but it can also be problematic on mobile or in cases when the behavior is unexpected.

Alias Popups When set, the tooltip presented by the GUI when hovering over an alias in a rule list only shows the alias description. When unset, the contents of the alias are included in the tooltip. For firewalls with large aliases, this may cause performance or browser rendering issues.

Disable Dragging When set, the GUI disables drag-and-drop on rule lists. Most users find drag-and-drop to be convenient and beneficial, thus the feature is enabled by default. Users who find the behavior undesirable can set this option.

Login Page Color Controls the color of the login page, which is independent of the theme.

Login Hostname When set, the GUI includes the hostname on the login form. This can be considered a security risk since it exposes information about the firewall to users who have not yet authenticated. If the firewall GUI is only reachable by authorized management clients, the convenience may outweigh the potential risk.

3.6 Advanced Configuration Options

System > Advanced contains numerous options of an advanced nature. These options customize the firewall behavior for more complex environments. Most administrators will not need to adjust these options for basic deployments.

Some of these options are covered in more detail in other sections of the documentation where their discussion is more topical or relevant, but they are all mentioned here with a brief description.

3.6.1 Admin Access Tab

The options on the Admin Access tab govern various methods for administering the firewall, including via the web interface, SSH, serial, and physical console.

webConfigurator (GUI)

Protocol

The protocol used by the GUI to accept web browser connections. May be one of:

HTTP Plain unencrypted HTTP. Insecure and basic, but widely compatible and less likely to have client issues. Should not be used in most cases, and should never be exposed to insecure networks.

HTTPS (SSL/TLS) Encrypted (“Secure”) HTTP. Protects communication between the client browser and the firewall GUI. Requires an SSL/TLS certificate to function. Depending on the browser and certificate configuration, there may be compatibility issues, but typically these are easily overcome by using current versions.

Note: The best practice is to use HTTPS so only encrypted traffic is exchanged between the GUI and clients.

SSL/TLS Certificate

The SSL/TLS Certificate to be used by the GUI in HTTPS (SSL/TLS) mode.

The firewall automatically generates a default self-signed certificate on the first boot. That is not an ideal situation, but is better than no encryption at all.

The primary disadvantage of a self-signed certificate is the lack of assurance of the identity of the host, since the certificate is not signed by a Certificate Authority trusted by the browser. Additionally, because for the bulk of Internet users such an invalid certificate should be considered a risk, modern browsers may restrict how such certificates are handled. Firefox, for example, gives a warning screen and forces the user to import the certificate and allow a permanent exception. Chrome shows a warning screen with a link to continue.

Tip: To use an externally signed SSL certificate and key, import them using the *Certificate Manager*, then select the certificate here.

Tip: The *ACME Package* can utilize the free Let’s Encrypt service to automatically obtain and update a signed certificate for the GUI or for other purposes on the firewall.

Tip: Another alternate technique is to generate a self-signed CA and then generate a GUI certificate from that CA. Export the CA from the firewall and then import that CA into client browsers manually. Using this method, all certificates signed by that CA will be trusted by browsers. Specifics vary by client platform.

Tip: To generate a new self-signed certificate for the GUI, connect using the console or ssh and from a shell prompt, run the following command:

```
# pfSsh.php playback generateguicert
```

TCP Port

The port used by the GUI for accepting connections from browsers. By default, the GUI uses HTTPS on port 443 with a redirect from port 80 for the best compatibility and ease of initial configuration. To change the port, enter a new port number into the TCP Port field.

Note: Moving the WebGUI to an alternate port is preferred by some administrators for security by obscurity reasons, though such practices should not be considered as offering any security benefit. Do not expose the GUI to untrusted networks such as the Internet.

Tip: Moving the GUI to another port will free up the standard web ports for use with port forwards or other services such as HAproxy.

Max Processes

The number of web server worker process used by the GUI when listening for client browser connections. The default value is 2.

If multiple administrators view the GUI at the same time and pages are taking too long to load, or are failing to load, then increase the Max Processes value.

WebGUI Redirect

Controls whether or not the firewall runs a redirect on port 80 so that if a browser attempts to access the firewall with HTTP, the firewall will accept the request and then redirect the browser to the TCP Port used by the GUI (e.g., HTTPS on port 443).

The redirect is enabled by default for ease of access and compatibility.

Disabling the redirect allows another daemon to bind to port 80.

HSTS

Controls whether the GUI web server sends the Strict-Transport-Security HTTPS response header (HSTS) to the browser. Check this box to disable the behavior.

HSTS forces the browser to use only HTTPS for future requests to the firewall FQDN to ensure it does not accidentally downgrade to an unencrypted connection.

Warning: When disabling HSTS, clients which visited the GUI when HSTS was enabled must perform browser specific steps for the change to take effect. Consult browser documentation for information on clearing cached HSTS behavior.

OCSP Must-Staple

Controls whether or not the GUI web server forcefully enables OCSP Stapling.

If the GUI SSL/TLS Certificate requires OCSP Stapling, this behavior is automatically enabled by the GUI web server. If the certificate property cannot be automatically determined by the firewall, this option can force the behavior.

Tip: Import the full CA and certificate chain or this option will be ignored by the GUI web server.

WebGUI Login Autocomplete

Controls whether or not the login form allows autocomplete so browsers can save the login credentials, for convenience.

In high-security environments, such as those that must adhere to specific security compliance standards, this behavior is not acceptable.

Note: This only controls autocomplete on the login form.

Warning: Few modern browsers respect this option. Many still offer to save passwords even when the form specifies that the browser must not allow the behavior. This behavior must be controlled or changed using browser options.

WebGUI login messages

Controls whether or not the firewall prints successful login messages to the console and system log.

On hardware with a PC speaker, these console messages generate a beep from the speaker, which some users find undesirable.

Checking this option stops the log message and the resulting beep.

Anti-lockout

Controls whether or not the firewall adds special rules to permit access to the WebGUI port and SSH port on the LAN interface by default.

These special rules override user-defined filter rules and prevent the user from accidentally locking themselves out of the firewall GUI or SSH. To control which LAN IP addresses may access the GUI and SSH using firewall rules, disable the anti-lockout rules.

When two or more interfaces are present, the firewall puts anti-lockout rules on the LAN interface; If only one interface is configured, the firewall places rules on that interface instead.

Warning: Filter rules must be in place to allow GUI access before enabling this option! If the LAN rules do not allow access to the GUI, removing the anti-lockout rule will block access to the GUI, potentially leaving the administrator without a means to reach the firewall.

Note: Resetting the LAN IP address from the console also resets the anti-lockout rule. If administrative access is unavailable after enabling this option, choose the console menu option 2, then choose to set the LAN IP address, and enter in the exact same IP address and accompanying information.

DNS Rebind Check

Controls whether or not the DNS resolver or forwarder performs DNS rebinding checks. These checks prevent the firewall from receiving DNS responses containing private IP addresses from DNS servers to prevent DNS rebinding attacks.

Note: When accessing the firewall by IP address, these checks are not enforced because the attack is only relevant when using a hostname.

Check this box to disable DNS rebinding protection if it interferes with GUI access or name resolution.

See also:

More detail on DNS rebinding attacks may be found on Wikipedia.

The most common case for disabling DNS rebinding checks is when the firewall is set to use an internal DNS server which will return private (RFC1918) answers for hostnames.

Tip: Instead of disabling all DNS rebinding protections, the checks can be selectively disabled on a per-domain basis in the DNS Resolver or DNS Forwarder. See *DNS Resolver and DNS Rebinding Protection* and *DNS Forwarder and DNS Rebinding Protection*.

Browser HTTP_REFERER enforcement

Controls whether or not the GUI checks and enforces HTTP_REFERER contents.

The GUI checks the referring URL sent by a client browser to ensure that the form was submitted from this firewall. This check prevents a form on another site from submitting a request to the firewall, changing an option when the administrator did not intend for that to happen.

This also breaks some convenience behaviors, such as having a page that links to various firewall devices, though the benefits of the check typically outweigh the advantage of those behaviors.

Alternate Hostnames

A list of Alternate Hostnames for the firewall allowed by DNS Rebind Checks and HTTP_REFERER Enforcement. To keep these features active, but alter their behavior slightly, add Alternate Hostnames.

By default, the GUI allows access to the hostname configured on the firewall and all IP addresses configured on the firewall. Hostnames in this field are allowed by the firewall for GUI access and for referring URL purposes.

Man-In-The-Middle Attack/Warning

If a browser attempts to access the GUI using an IP address that is not configured on the firewall, such as a port forward from another firewall, the GUI prints a message indicating that access to the firewall may be compromised due to a Man-In-The-Middle (MITM) attack.

If such a forwarding was deliberatey configured on this firewall or on a firewall upstream, the message may be safely ignored. If access to the firewall should have been direct, then take great care before logging in to ensure the login credentials are not being routed through an untrusted system.

Access is not disabled by the firewall in this case, it only prints a warning, so there is no option to disable this behavior.

Browser Tab Text

By default, the GUI prints the firewall hostname first in the page/tab title, followed by the page name. To reverse this behavior and show the page name first and hostname second, check Display page name first in browser tab.

Administrators who access many firewalls at the same time in separate tabs tend to prefer having the hostname first (default). Administrators who access one firewall with many pages in separate tabs tend to prefer having the page name first.

Secure Shell (SSH)

The Secure Shell (SSH) server provides remote console access and file management. A user can connect with any standard SSH client, such as the OpenSSH command line ssh client, PuTTY, SecureCRT, or iTerm.

When using SSH, both the admin username and root username are accessible using the admin account credentials.

Users in the User Manager that have the User - System - Shell account access privilege is also allowed to login over ssh. These users do not have root access privileges, and do not print the menu when they login because many of the options require root privileges.

Tip: To grant users additional shell privileges, use the *sudo package*.

File transfers to and from the firewall are also possible by using a Secure Copy (SCP) client such as the OpenSSH command line scp, FileZilla, WinSCP or Fugu. To use SCP, connect as the root user, not admin. If a custom user has the User - System - Copy files permission, or all access, then they may also utilize SCP.

Tip: SSH clients must be kept up-to-date. As time goes on, security standards evolve and the SSH server settings utilized by SSH servers will change. Outdated clients may not be able to connect using the strong security keys and algorithms required by sshd. If a client will not connect, check for an update from the vendor.

Enable Secure Shell

To enable the SSH daemon, check Enable Secure Shell. After saving with this option enabled, the firewall will generate SSH keys if they are not already present and then start the SSH daemon.

SSHd Key Only

This option controls which authentication methods the SSH daemon allows for clients. It can be set to one of the following values:

Password or Public Key Allows a user to authenticate with either a valid password or valid key. This is the default behavior.

Public Key Only Restricts authentication to only valid keys, passwords are not allowed.

Require Both Password and Public Key Requires a valid password and a valid key.

Key-based logins are a much more secure practice, though it does take more preparation to configure.

Add user keys for key-based login by editing users in the User Manager (*User Management and Authentication*). When editing a user, paste the allowed public keys into the Authorized Keys text field for the account.

Allow Agent Forwarding

Controls whether or not the SSH daemon allows agent forwarding for clients.

Agent forwarding allows a user to run an SSH agent on their client system and connect to the firewall, and then to other remote SSH servers using the key from their agent. In this case, the user does not need to have their private keys on the firewall but can still use key-based authentication to remote servers.

Use of an SSH agent can be considered a security issue in certain cases. Additionally, the firewall is not intended to be a general purpose SSH client or intermediate system, thus this feature is disabled by default.

SSH Port

Controls the port used by the SSH daemon to accept client connections. To change the port, type the new port into the SSH Port box.

Moving the SSH server to an alternate port provides a negligible security improvement, and frees up the port for other uses.

Tip: Brute force SSH scanners focus on hitting TCP port 22 but if the daemon is open to the Internet on another port, it will eventually be found and hit by scanners.

Best Practices for SSH

If this firewall is installed in an environment that requires leaving SSH access unrestricted by firewall rules, which is dangerous, we strongly recommended the following actions:

Change the SSH Port Moving to a random alternate port prevents log noise from many, but not all, brute-force SSH login attempts and casual scans. It can still be found with a port scan, however.

Force Key-Based Authentication Key-based authentication must always be used by publicly accessible SSH servers to eliminate the possibility of successful brute force attacks. Set SSHd Key Only to either *Public Key Only* or *Require Both Password and Public Key*.

Multiple unsuccessful logins from the same IP address will result in locking out the IP address trying to authenticate, but that alone is not considered sufficient protection.

Login Protection

The sshguard daemon is used by the firewall to protect against brute force logins for both the GUI and SSH connections. The options in this section fine-tune the behavior of this protection.

Threshold The total score value above which sshguard will block clients. Most attacks have a score of 10, the default threshold value is 30.

Blocktime the initial minimum number of seconds to block attackers who have exceeded the Threshold value. The default value is 120 seconds. Repeat offenders are blocked for increasingly longer amounts of time (1.5x for each repetition).

Note: Attackers are unblocked at random intervals so actual block time will be longer than stated. This prevents clients from predicting the timing to optimize targeted attacks.

Detection Time The amount of time, in seconds, attackers are remembered by sshguard since their last offense before it resets their score. Default is 1800 seconds.

Whitelist A list of subnets which are excluded from login protection. This lowers security but is generally acceptable from specific secure management networks.

For example, it may be necessary to add entries for network monitoring systems which probe the SSH port but do not login. Otherwise, such systems may be flagged as attackers.

Serial Communications

If the firewall is running on hardware without a monitor or if it will be running “headless” (without keyboard and video attached), then the serial console can be enabled to maintain physical control, so long as the hardware has a serial port (not USB).

If hardware is detected which has no VGA port, the serial console is forced on and cannot be disabled, and the serial options are all hidden except for the speed.

Serial Terminal

When Serial Terminal is set, the operating system enables the console on the first serial port. This console will receive kernel boot messages and a menu after the firewall has finished booting. This will not disable the onboard keyboard and video console.

To connect to the serial console, use a null modem cable connected to a serial port or adapter on another PC or serial device.

See also:

For more information on connecting to a serial console, see *Connecting to a Serial Console* and *Start a Serial Client*.

When making any changes to the serial console, the firewall must be rebooted before they take effect.

Serial Console Speed

The default serial console speed is *115200* bps and almost all hardware works well at that speed. In rare cases, a slower speed may be required which can be set here by picking the desired speed from the Serial Speed drop-down.

When upgrading from an older version, this may remain at an older value such as *9600* or *38400* to match the BIOS on older hardware. Increasing the speed to *115200* is almost always safe and more useful than slower speeds.

Primary Console

On hardware with both the serial console enabled and a VGA port available, the Primary Console selector chooses which is the preferred console, so it will receive the boot log messages. Other OS kernel messages will show up on all console connections, and both consoles will have a usable menu.

In cases where the boot cannot complete, the preferred console must be used to resolve the problem, such as reassigning interfaces.

Console Menu

Normally the firewall always presents the menu on the console, and the menu will be available as long as someone has physical access to the console. In high-security environments this is not desirable.

This option adds password protection to the console. The console accepts the same usernames and passwords used to access the GUI. After setting this option, the firewall must be rebooted before it takes effect.

3.6.2 Firewall/NAT Tab

Firewall Advanced

IP Do-Not-Fragment compatibility

This option is a workaround for operating systems which generate fragmented packets with the "don't fragment" (DF) bit set. Linux NFS (Network File System) is known to do this, as well as some VoIP implementations.

When this option is enabled, the firewall will not drop these malformed packets but instead it will clear the DF bit. The firewall will also randomize the IP identification field of outgoing packets to compensate for operating systems that set the DF bit but set a zero IP identification header field.

IP Random ID generation

If Insert a stronger ID into IP header of packets passing through the filter is checked, the firewall replaces the IP identification field of packets with random values to compensate for operating systems that use predictable values. This option only applies to packets that are not fragmented after the optional packet reassembly.

Firewall Optimization Options

The optimization mode controls how the firewall expires state table entries:

Normally The standard optimization algorithm, which is optimal for most environments.

High Latency Used for high latency links, such as satellite links. Expires idle connections later than default.

Aggressive Expires idle connections quicker. More efficient use of CPU and memory but can drop legitimate connections earlier than expected. This option can also improve performance in high traffic deployments with lots of connections, such as web services.

Conservative Tries to avoid dropping any legitimate connections at the expense of increased memory usage and CPU utilization. Can aid in environments that require long-lived but mostly idle UDP connections, such as VoIP.

The table *Firewall Optimization Details* contains the values chosen by PF for each optimization algorithm. The values are taken from the PF source code. The first line is the raw value, second line is human readable:

Table 1: Firewall Optimization Details

	Normal	High Latency	Conservative	Aggressive
tcp.first First TCP packet	60 1min	180 3min	3600 60min	30 30sec
tcp.opening No response yet	30 30sec	35 35sec	900 15min	5 5sec

tcp.established Established	86400 24h	86400 24h	432000 5days	18000 5h
tcp.closing Half closed	900 15min	905 15min + 5sec	3600 1h	60 60sec
tcp.finwait Got both FINs	45 45sec	50 50sec	600 10min	30 30sec
tcp.closed Got an RST	90 90sec	95 95sec	180 3min	30 30sec
tcp.tsdiff Allowed TS diff	30 30sec	60 60sec	60 60sec	10 10sec

Disable Firewall

When Disable all packet filtering is set, the firewall becomes a routing-only platform. This is accomplished by disabling pf entirely, and as a consequence, NAT is disabled since it is also handled by pf.

Tip: To disable *only* NAT, do not use this option. Consult *Disabling Outbound NAT* for more information on controlling outbound NAT behavior.

Disable Firewall Scrub

When set, the scrubbing option in pf is disabled. The scrub action in pf can interfere with NFS, and in rare cases, with VoIP traffic as well. By default, the firewall uses the fragment reassemble option which reassembles fragmented packets before sending them on to their destination, when possible. More information on the scrub feature of pf can be found in the OpenBSD PF Scrub Documentation.

Note: Disabling scrub also disables other features that rely on scrub to function, such as DF bit clearing and ID randomization. Disabling scrub does not disable MSS clamping if it is active for VPNs, or when an MSS value is configured on an interface.

Firewall Adaptive Timeouts

Adaptive Timeouts control state handling in pf when the state table is nearly full. Using these timeouts, a firewall administrator can control how states are expired or purged when there is little or no space remaining to store new connection states.

Adaptive Timeouts are enabled by default and the default values are calculated automatically based on the configured Firewall Maximum States value.

Adaptive Start Adaptive scaling is started once the state table reaches this level, expressed as a number of states. Adaptive Start defaults to 60% of Firewall Maximum States.

Adaptive End When the size of the state table reaches this value, expressed as a number of state table entries, all timeout values are assumed to be zero, which causes pf to purge all state entries immediately. This setting defines the scale factor, it should be set greater than the total number of states allowed. Adaptive End defaults to 120% of Firewall Maximum States.

When the number of connection states exceeds the threshold set by Adaptive Start, timeout values are scaled linearly with factor based on the number of states used between the Start and End state counts. The timeout adjustment factor is calculated as follows: (Number of states until the Adaptive End value is reached) / (Difference between the Adaptive End and Adaptive Start values).

Note: As an example, consider a firewall with Adaptive Start set to 600000, Adaptive End set to 1200000 and Firewall Maximum States set to 1000000. In this situation, when the state table size reaches 900000 entries the state timeouts will be scaled to 50% of their normal values.

(1,200,000 - 900,000) / (1,200,000 - 600,000) = 300,000 / 600,000 = 0.50, 50%

Continuing the example, when the state table is full at 1,000,000 states the timeout values will be reduced to 1/3 of their original values.

Tip: The state table usage indicator on the dashboard will change color and text when the state table size crosses these thresholds.

Firewall Maximum States

This value is the maximum number of connections the firewall can hold in its state table. The default size is calculated based on 10% of total RAM. This default value is sufficient for most installations, but can be adjusted higher or lower depending on the load and available memory.

Each state consumes approximately 1 KB of RAM, or roughly 1 MB of RAM for every 1000 states. The firewall must have adequate free RAM to contain the entire state table before increasing this value. Firewall states are discussed further in *Stateful Filtering*.

Tip: On a firewall with 8GB of RAM the state table would have a default size of approximately 800,000 states. A custom Firewall Maximum States value of 4,000,000 would consume about 4GB of RAM, half the available 8GB total.

Firewall Maximum Table Entries

This value defines the maximum number of entries that can exist inside of address tables used by the firewall for collections of addresses such as aliases, ssh/GUI lockout records, hosts blocked by snort alerts, and so on. By default this is 200,000 entries. If the firewall has features enabled which can load large blocks of address space into aliases such as URL Table aliases or the pfBlockerNG package, then increase this value to comfortably include at least double the total amount of entries contained in all aliases combined.

Firewall Maximum Fragment Entries

When scrub is enabled the firewall maintains a table of packet fragments waiting to be reassembled. By default this table can hold 5000 fragments. In rare cases a network may have an unusually high rate of fragmented packets which can require more space in this table. When this limit is reached, the following log message will appear in the main system log:

```
kernel: [zone: pf frag entries] PF frag entries limit reached
```

Static Route Filtering

The Bypass firewall rules for traffic on the same interface option applies if the firewall has one or more static routes defined. If this option is enabled, traffic that enters and leaves through the same interface will not be checked by the firewall. This may be required in situations where multiple subnets are connected to the same interface, to avoid blocking traffic that is passed through the firewall in one direction only due to asymmetric routing. See *Bypass Firewall Rules for Traffic on Same Interface* for a more in-depth discussion on that topic.

Disable Auto-added VPN rules

By default, when IPsec is enabled firewall rules are automatically added to the appropriate interface which will allow the tunnel to establish. When Disable Auto-added VPN rules is checked, the firewall will not automatically add these rules. By disabling these automatic rules, the firewall administrator has control over which addresses are allowed to connect to a VPN. Further information on these rules can be found at *VPNs and Firewall Rules*.

Disable Reply-To

In a Multi-WAN configuration the firewall has a beneficial default behavior that ensures traffic leaves the same interface it arrived through. This is accomplished using the pf keyword reply-to which is added automatically to interface tab firewall rules for WAN-type interfaces. When a connection matches a rule with reply-to, the firewall remembers the path through which the connection was made and routes the reply traffic back to the gateway for that interface.

Tip: WAN-type interfaces are interfaces which have a gateway set on their Interfaces menu entry configuration, or interfaces which have a dynamic gateway such as DHCP, PPPoE, or assigned OpenVPN, GIF, or GRE interfaces.

In situations such as bridging, this behavior is undesirable if the WAN gateway IP address is different from the gateway IP address of the hosts behind the bridged interface. Disabling reply-to will allow clients to communicate with the proper gateway.

Another case that has issues with reply-to involves static routing to other systems in a larger WAN subnet. Disabling reply-to in this case would help ensure that replies return to the proper router instead of being routed back to the gateway.

Allow APIPA

Automatic Private IP Addressing (APIPA), or IPv4 Link-Local addressing, uses a special subnet of 169.254.0.0/ 16. This traffic is for local links only (same L2), it must not be routed or traverse a firewall. As such, inbound traffic from these addresses is automatically blocked by internal firewall rules by default.

When Allow APIPA traffic is checked, the default block rules are removed, and user firewall rules can control the traffic.

There are some use cases which utilize these addresses for private communication on an interface, such as AWS VPC BGP, and in those cases, the option can be enabled along with carefully crafted manual firewall rules.

Warning: When this option is enabled, take care to never allow APIPA traffic to match policy routing rules. If APIPA traffic matches policy routing rules, behavior is unpredictable. There have been reports of such errors leading to packet loops and unexpectedly high resource usage. See Redmine Issue #2073 for more.

Aliases Hostnames Resolve Interval

This option controls how often hostnames in aliases are resolved and updated by the filterdns daemon. By default this is 300 seconds (5 minutes). In configurations with a small number of hostnames or a fast/low-load DNS server, decrease this value to pick up changes faster.

Check Certificate of Alias URLs

When Verify HTTPS certificates when downloading alias URLs is set, the firewall will require a valid HTTPS certificate for web servers used in URL table aliases. This behavior is more secure, but if the web server is private and uses a self-signed certificate, it can be more convenient to ignore the validity of the certificate and allow the data to be downloaded.

Warning: The best practice is to always use a server certificate with a valid chain of trust for this type of role, rather than weakening security by allowing a self-signed certificate.

Bogon Networks

The Update Frequency drop-down for Bogon Networks controls how often these lists are updated. Further information on bogon networks may be found in *Block Bogon Networks*.

Network Address Translation

NAT Reflection for Port Forwards

The NAT Reflection mode for port forwards option controls how NAT reflection is handled by the firewall. These NAT redirect rules allow clients to access port forwards using the public IP addresses on the firewall from within local internal networks.

See also:

Refer to *NAT Reflection* for a discussion on the merits of NAT Reflection when compared to other techniques such as Split DNS.

There are three possible modes for NAT Reflection:

Disabled The default value. When disabled, port forwards are only accessible from WAN and not from inside local networks.

Pure NAT This mode uses a set of NAT rules to direct packets to the target of the port forward. It has better scalability, but it must be possible to accurately determine the interface and gateway IP address used for communication with the target at the time the rules are loaded. There are no inherent limits to the number of ports other than the limits of the protocols. All protocols available for port forwards are supported.

When this option is enabled, Automatic Outbound NAT for Reflection must also be enabled if the clients and servers are in the same local network.

NAT + Proxy *NAT + proxy* mode uses a helper program to send packets to the target of the port forward. The connection is received by the reflection daemon and it acts as a proxy, creating a new connection to the local server. This behavior puts a larger burden on the firewall, but is useful in setups where the interface and/or gateway IP address used for communication with the target cannot be accurately determined at the time the rules are loaded. *NAT + Proxy* reflection rules are not created for ranges larger than 500 ports and will not be used for more than 1000 ports total between all port forwards. This feature only supports TCP port forwards.

Individual NAT rules have the option to override the global NAT reflection configuration, so they may have NAT reflection forced on or off on a case-by-case basis.

Reflection Timeout

The Reflection Timeout setting forces a timeout on connections made when performing NAT reflection for port forwards in *NAT + Proxy* mode. If connections are staying open and consuming resources, this option can mitigate that issue.

NAT Reflection for 1:1 NAT

When checked, this option adds additional reflection rules which enable access to 1:1 mappings of external IP addresses from internal networks. This gives the same functionality that already exists for port forwards, but for 1:1 NAT. There are complex routing scenarios that may render this option ineffective.

This option only affects the *inbound* path for 1:1 NAT, not outbound. The underlying rule style is similar to the *Pure NAT* mode for port forwards. As with port forwards, there are per-entry options to override this behavior.

Automatic Outbound NAT for Reflection

When checked, this option automatically creates outbound NAT rules which assist reflection rules that direct traffic back out to the same subnet from which it originated. These additional rules allow Pure NAT and 1:1 NAT Reflection to function fully when the clients and servers are in the same subnet. In most cases, this box must be checked for NAT Reflection to work.

Note: This behavior is necessary because when clients and servers are in the same subnet, the traffic source must be changed so that the connection appears to originate from the firewall. Otherwise, the return traffic will bypass the firewall and the connection will not succeed.

TFTP Proxy

The built-in TFTP proxy will proxy connections to TFTP servers outside the firewall, so that client connections may be made to remote TFTP servers. Ctrl-click or shift-click to select multiple entries from the list. If no interfaces are chosen, the TFTP proxy service is deactivated.

State Timeouts

The State Timeout section allows fine-tuning of the state timeouts for various protocols. These are typically handled automatically by the firewall and the values are dictated by the *Firewall Optimization Options*. In rare cases, these timeouts may need adjusted up or down to account for irregularities in device behavior or site-specific needs.

All of the values are expressed in *seconds*, and control how long a connection in that state will be retained in the state table.

See also:

Descriptions in the following options reference firewall state conditions as described in *Interpreting States*.

TCP First The first packet of a TCP connection.

TCP Opening The state before the destination host has replied (e.g. SYN_SENT:CLOSED).

TCP Established An established TCP connection where the three-way handshake has been completed.

TCP Closing One side has sent a TCP FIN packet.

TCP FIN Wait Both sides have exchanged FIN packets and the connection is shutting down. Some servers may continue to send packets during this time.

TCP Closed One side has sent a connection reset (TCP RST) packet.

TCP Tsdiff The allowed TCP timestamp difference.

UDP First The first UDP packet of a connection has been received.

UDP Single The source host has sent a single packet but the destination has not replied (e.g. SINGLE:NO_TRAFFIC).

UDP Multiple Both sides have sent packets.

ICMP First an ICMP packet has been received.

ICMP Error an ICMP error was received in response to an ICMP packet.

Other First, Other Single, Other Multiple The same as UDP, but for other protocols.

3.6.3 Networking Tab

IPv6 Options Allow

IPv6

The Allow IPv6 option controls a set of block rules which prevent IPv6 traffic from being handled by the firewall.

Note: This option does not disable IPv6 functions or prevent it from being configured, it only controls traffic flow.

When the option is enabled, IPv6 traffic will be allowed when permitted by firewall rules and/or automatic rules, depending on the firewall configuration. This option is enabled by default on new configurations.

When the option is unchecked, all IPv6 traffic will be blocked. This behavior is similar to how IPv6 was treated before it was supported by AZTCO-FW® software. Configurations imported from or upgraded from versions older than 2.1 will have this option unchecked, so they behave consistently after upgrade.

IPv6 over IPv4 Tunneling

The Enable IPv6 over IPv4 Tunneling option enables forwarding for IP protocol 41/RFC 2893 to an IPv4 address specified in the IPv4 address of Tunnel Peer field.

When configured, this forwards all incoming protocol 41/IPv6 traffic to a host behind this firewall instead of handling it locally.

Tip: Enabling this option does not add firewall rules to allow the protocol 41 traffic. A rule must exist on the WAN interface to allow the traffic to pass through to the local receiving host.

Prefer IPv4 over IPv6

When set, this option causes the firewall *itself* to prefer sending traffic to IPv4 hosts instead of IPv6 hosts when a DNS query returns results for both.

In rare cases when the firewall has partially configured, but not fully routed, IPv6 this can allow the firewall to continue reaching Internet hosts over IPv4.

Note: This option controls the behavior of the firewall itself, such as when polling for updates, package installations, downloading rules, and fetching other data. It cannot influence the behavior of clients behind the firewall.

IPv6 DNS Entry

This option controls whether or not the firewall creates local DNS entries for the firewall itself with IPv6 addresses, when available.

By default (unchecked), the firewall automatically adds DNS entries for itself using its local IPv4 and IPv6 interface addresses. In some cases, such as with dynamic IPv6 addresses like tracked interfaces, the IPv6 address may disappear or change and clients may attempt to use an outdated address until their cached DNS response expires.

When the option is checked, the firewall only adds DNS entries for its IPv4 addresses.

DHCP6 DUID

This option controls the DHCPv6 Unique Identifier (DUID) used by the firewall when requesting an IPv6 address. The firewall generates a DUID automatically, but in some cases, an administrator may want to use a different DUID. For example, if the operating system was reinstalled and the firewall should use the same DUID it had in the past, or if an upstream network administrator requires a specific DUID.

Note: Most users do not need to change this to any specific value, the default behavior is fine for nearly all environments. When in doubt, leave it alone unless directed to change it by an upstream network provider.

There are several possible DUID formats that this option can accept, chosen by the drop-down menu. When a format is chosen, the GUI displays a different set of input boxes specific to the selected format. The exact format depends upon the needs of the network administrator (e.g. ISP, datacenter, etc) and they would provide the format and values.

The available DUID formats are:

Raw DUID represented exactly as observed in a DUID file or in logs. Entered as:

Raw DUID A single text area in which the DUID can be entered.

This option also includes a ⧉ Copy DUID button which copies the DUID from the placeholder (automatically generated by the firewall) into the text box so that the existing DUID can easily be placed into the configuration.

DUID-LLT DUID format with Link-Layer Address Plus Time. Entered as:

Time (in seconds) since January 1st, 2000 UTC

Link-Layer Address The link-layer address (MAC) of an interface on the firewall in the format xx:xx:xx:xx:xx:xx.

DUID-EN DUID assigned by a vendor based on Enterprise Number. Entered as:

Enterprise Number IANA Private Enterprise Number of the vendor.

Identifier Variable length identifier in the format xx:xx:xx:xx. The length depends upon the vendor.

DUID-LL DUID based on only Link-Layer Address. Entered as:

Link-Layer Address The link-layer address (MAC) of an interface on the firewall in the format xx:xx:xx:xx:xx:xx.

DUID-UUID DUID based on the host Universally Unique Identifier (UUID). Entered as:

DUID-UUID The UUID for this host in the format nnnnnnnn-nnnn-nnnn-nnnn-nnnnnnnnnnnn

Network Interfaces

Hardware Checksum Offloading

When checked, this option disables hardware checksum offloading on the network cards. Checksum offloading is usually beneficial as it allows the checksum to be calculated (outgoing) or verified (incoming) in hardware at a much faster rate than it could be handled in software.

Note: When checksum offloading is enabled, a packet capture will see empty (all zero) or flag incorrect packet checksums. These are normal when checksum handling is happening in hardware.

Checksum offloading is broken in some hardware, particularly Realtek cards and virtualized/emulated cards such as those on Xen/KVM. Typical symptoms of broken checksum offloading include corrupted packets and poor throughput performance.

Tip: In virtualization cases such as Xen/KVM it may be necessary to disable checksum offloading on the host as well as the VM. If performance is still poor or has errors on these types of VMs, switch the type of NIC if possible.

Hardware TCP Segmentation Offloading

Checking this option will disable hardware TCP segmentation offloading (TSO, TSO4, TSO6). TSO causes the NIC to handle splitting up packets into MTU-sized chunks rather than handling that at the OS level. This can be faster for servers and appliances as it allows the OS to offload that task to dedicated hardware, but when acting as a firewall or router this behavior is highly undesirable as it actually increases the load as this task has already been performed elsewhere on the network, thus breaking the end-to-end principle by modifying packets that did not originate on this host.

Warning: This option is not desirable for routers and firewalls, but can benefit workstations and appliances. It is disabled by default, and should remain disabled unless the firewall is acting primarily or solely in an appliance/endpoint role.

Do not uncheck this option unless directed to do so by a support representative. This offloading is broken in some hardware drivers, and can negatively impact performance on affected network cards and roles.

Hardware Large Receive Offloading

Checking this option will disable hardware large receive offloading (LRO). LRO is similar to TSO, but for the incoming path rather than outgoing. It allows the NIC to receive a large number of smaller packets before passing them up to the operating system as a larger chunk. This can be faster for servers and appliances as it offloads what would normally be a processing-heavy task to the network card. When acting as a firewall or router this is highly undesirable as it delays the reception and forwarding of packets that are not destined for this host, and they will have to be split back up again on the outbound path, increasing the workload significantly and breaking the end-to-end principle.

> Warning: This option is not desirable for routers and firewalls, but can benefit workstations and appliances. It is disabled by default, and should remain disabled unless the firewall is acting primarily or solely in an appliance/endpoint role.
>
> Do not uncheck this option unless directed to do so by a support representative. This offloading is broken in some hardware drivers, and can negatively impact performance on affected network cards and roles.

hn ALTQ Support

Checking this option will enable support for ALTQ traffic shaping on hn(4) network interfaces in Hyper-V.

For ALTQ to work on hn(4) interfaces, the operating system must disable the multi-queue API which may reduce the system capability to handle traffic. The administrator must decide if this reduction in performance is worth the benefit of traffic shaping.

The firewall must be rebooted for this setting to take effect.

Suppress ARP messages

The firewall makes a log entry in the main system log when an IP address appears to switch to a different MAC address. This log entry notes that the device has moved addresses, and records the IP address and the old and new MAC addresses.

This event can be completely benign behavior (e.g. NIC teaming on a Microsoft server, a device being replaced) or a legitimate client problem (e.g. IP conflict), and it could show up constantly or rarely if ever. It all depends on the network environment.

The best practice is to allow these ARP messages to be printed to log since there is a chance it will report a problem worth the attention of a network administrator. However, if the network environment contains systems which generate these messages while operating normally, suppressing the errors can make the system log more useful as it will not be cluttered with unneeded log messages.

Reset All States

When set, if an interface IP address changes, the firewall will reset the entire state table instead of only clearing states for the old interface IP address.

This behavior is potentially disruptive, and is off by default. In single WAN environments, this is not typically any more disruptive than the WAN address changing, since clients already have to reestablish all connections.

In most cases, this behavior is not necessary, but it can help in certain situations where WAN addresses change rapidly and the normal behavior misses states for former IP addresses.

3.6.4 Miscellaneous Tab

Proxy Support

If this firewall resides in a network which requires a proxy for outbound Internet access, enter the proxy options in this section so that requests from the firewall for items such as packages and updates will be sent through the proxy.

Proxy URL This option specifies the location of the proxy for making outside connections. It must be an IP address or a fully qualified domain name.

Proxy Port The port to use when connecting to the proxy URL. By default the port is 8080 for HTTP proxy URLs, and 443 for SSL proxy URLs. The port is determined by the proxy, and may be a different value entirely (e.g. 3128). Check with the proxy administrator to find the proper port value.

Proxy Username If required, this is the username that is sent for proxy authentication.

Proxy Password If required, this is the password associated with the username set in the previous option.

Load Balancing

When AZTCO-FW® software is directed to perform load balancing, successive connections will be redirected in a round robin manner to a gateway, balancing the load across all available paths. The options in this section alter or fine-tune that behavior.

Sticky Connections When active, connections from the same source are sent through the same gateway, rather than being sent in a purely round-robin manner.

This "sticky" association will exist as long as states are in the table for connections from a given source address (e.g. the IP address of a user). Once the states for that source expire, so will the sticky association. Further connections from that source host will be redirected to the next available gateway in the group.

This behavior can help with protocols such as HTTPS and FTP, where the server may be strict about all connections coming from the same source. The downside of this behavior is that balancing is not as efficient, a heavy user could dominate a single WAN rather than having their connections spread out.

Source Tracking Timeout Controls how long the sticky association will be maintained for a host after the all of the states from that host expire. The value is specified in seconds.

By default, this value is not set, so the association is removed as soon as the states expire. If sticky connections appear to work initially but seem to stop partway through sessions, increase this value to hold an association longer. Web browsers often hold open connections for a while as users are on a site, but if there is a lot of idle time, connections may be closed and states may expire.

Power Savings

When Enable PowerD is checked, the powerd daemon is started. This daemon monitors the system and can lower or raise the CPU frequency based on system activity. If processes need the power, the CPU speed will be increased as needed. This option will lower the amount of heat a CPU generates, and may also lower power consumption.

Note: The behavior of this option depends greatly on the hardware in use. In some cases, the CPU frequency may lower but have no measurable effect on power consumption and/or heat, where others will cool down and use considerably less power. It is considered safe to run, but is left off by default unless supported hardware is detected.

The mode for powerd may also be selected for three system states:

AC Power Normal operation connected to AC power.

Battery Power Mode to use when the firewall is running on battery. Support for battery power detection varies by hardware.

Unknown Power Mode used when powerd cannot determine the power source.

Four modes choices exist for each of these states:

Maximum Keeps the performance as high as possible at all times.

Minimum Keeps performance at its lowest, to reduce power consumption.

Adaptive Tries to balance savings by decreasing performance when the system is idle and increasing when busy.

Hiadaptive Similar to adaptive but tuned to keep performance high at the cost of increased power consumption. It raises the CPU frequency faster and drops it slower. This is the default mode.

Note: Some hardware requires powerd running to operate at its maximum attainable CPU frequency. If the firewall device does not have powerd enabled but always runs at what appears to be a low CPU frequency, enable powerd and set it to *Maximum* for at least the AC Power state.

Watchdog

Certain firewall hardware includes a Watchdog feature which can reset the hardware when the watchdog daemon can no longer interface with the hardware after a specified timeout. This can increase reliability by resetting a unit when a hard lock is encountered that might otherwise require manual intervention.

The downside to any hardware watchdog is that any sufficiently busy system may be indistinguishable from one that has suffered a hard lock.

Enable Watchdog When checked, the watchdogd daemon is run which attempts to latch onto a supported hardware watchdog device.

Watchdog Timeout The time, in seconds, after which the device will be reset if it fails to respond to a watchdog request. If a firewall regularly has a high load and triggers the watchdog accidentally, increase the timeout.

Kernel Page Table Isolation (PTI)

Kernel PTI is a method for working around CPU vulnerabilities such as Meltdown. By exploiting that vulnerability without Kernel PTI, kernel memory could be accessed by unprivileged users on affected CPUs.

Note: While more secure, this protection can incur a performance penalty. If untrusted users do not have access to run arbitrary code on the firewall, it can be disabled without significant security risk.

Kernel PTI is active by default only on CPUs affected by the vulnerability.

This option forces the workaround off, and requires a reboot to change.

If a vulnerable CPU is not detected, PTI is disabled by default and this option will have no effect.

The current state of Kernel PTI is printed below the option.

Microarchitectural Data Sampling (MDS) Mitigation

Microarchitectural Data Sampling (MDS) mitigation is a method for working around weaknesses in Intel CPUs which support hyperthreading. By exploiting MDS without mitigation in place, kernel memory could be accessed by unprivileged users on affected CPUs.

Note: While more secure, this protection can incur a performance penalty. If untrusted users do not have access to run arbitrary code on the firewall, it can be disabled without significant security risk.

This option controls which method of MDS mitigation is used, if any. Changing the option requires a reboot to activate. The following modes are available:

Default The default operating system behavior. As of this writing, the default behavior is to disable MDS mitigation.

Mitigation Disabled Forcefully disable MDS mitigation.

VERW instruction (microcode) mitigation enabled Use VERW instruction mitigation, implemented in CPU microcode, to mitigate MDS. This is the fastest and most optimal way to mitigate MDS, but it requires support in the CPU microcode for this instruction.

Software sequence mitigation enabled Mitigates MDS by using software sequences, which is much slower, but safer.

Automatic VERW or Software selection When set to Automatic, the operating system will attempt to use VERW instructions if they are available and software in all other cases.

The current state of MDS mitigation is printed below the option.

Schedules

The Do not kill connections when schedule expires option controls whether or not states are cleared when a scheduled rule transitions into a state that would block traffic. If unchecked, connections are terminated when the schedule time has expired. If checked, connections are left alone and will not be automatically closed by the firewall.

Gateway Monitoring

State Killing on Gateway Failure

When using Multi-WAN, by default the monitoring process will not flush states when a gateway goes into a down state. Flushing states for each gateway event can be disruptive in situations where a gateway is unstable.

The Flush all states when a gateway goes down option overrides the default behavior, clearing states for all existing connections when any gateway fails. Clearing states can help redirect traffic for long-lived connections such as VoIP phone/trunk registrations to another WAN, but it can also disrupt ongoing connections if a lesser-used gateway is flapping which would still kill all states when it fails.

More information on how this impacts Multi-WAN can be found in *State Killing/Forced Switch*.

Warning: When this is triggered, the entire state table is cleared. This is necessary because it is not possible to kill all states for the failing WAN and the LAN-side states associated with the failing WAN. Removing states on the WAN side alone is ineffective, the LAN-side states must be cleared as well.

Skip Rules When Gateway is Down

By default, when a rule has a specific gateway set and this gateway is down, the gateway is omitted from the rule and traffic is sent via the default gateway.

The Do not create rules when gateway is down option overrides that behavior and the entire rule is omitted from the ruleset when the gateway is down. Instead of flowing via the default gateway, the traffic will match a different rule instead. This is useful if traffic must only ever use one specific WAN and never flow over any other WAN.

Tip: When utilizing this option, create a reject or block rule underneath the policy routing rule with the same matching criteria. This will prevent the traffic from potentially matching other rules below it in the ruleset and taking an unintended path.

RAM Disk Settings

The /tmp and /var directories are used for writing files and holding data that is temporary and/or volatile. Using a RAM disk can reduce the amount of writing that happens on disks in the firewall. Modern SSDs do not have disk write concerns as older drives once did, but it can still be a concern when running from lower quality flash storage such as USB thumb drives.

This behavior has the benefit of keeping most of the writes off of the disk in the base system, but packages may yet write frequently to the drive. It also requires additional handling to ensure data such as RRD graphs and DHCP leases are retained across reboots. Data for both is saved during a proper shutdown or reboot, and also periodically if configured.

Use RAM Disks When checked, a memory disk is created at boot time for /tmp and /var/ and the associated structure is initialized. When this setting is toggled, a reboot is required and forced on save.

Warning: The size of RAM disks is limited by the amount of available kernel memory. The actual limit is calculated and printed in the GUI underneath the size options.

/tmp RAM Disk Size The size of the /tmp RAM disk, in MiB. The default value is 40, but should be set higher if there is available RAM and kernel memory.

/var RAM Disk Size The size of the /var RAM disk, in MiB. The default value is 60, but should be set much higher, especially if packages will be used. 512-1024 is a better starting point, depending on the available firewall RAM and kernel memory.

Periodic RAM Disk Data Backups These options control how frequently data in RAM disks is backed up. If the firewall is rebooted unexpectedly, the last backup is restored when the firewall boots. The lower the value, the less data that will be lost in such an event, but more frequent backups write more to the disk.

RRD Data The time, in hours, between periodic backups of RRD files containing graph data.

DHCP Leases The time, in hours, betweenperiodic backups of the DHCP lease databases.

Log Directory The time, in hours, between periodic backups of the system log directory.

Warning: Aside from the points mentioned above, there are several items to be cautious about when choosing whether or not to use the RAM disk option. Used improperly, this option can lead to data loss or other unexpected failures.

Utilize remote syslog to send the logs to another device on the network rather than risking losing data from unexpected outages.

Packages may not properly account for the use of RAM disks, and may not function properly at boot time or in other ways. Test each package, including whether or not it works immediately after a reboot.

These are RAM disks, so the amount of RAM available to other programs will be reduced by the amount of space used by the RAM disks. For example if the firewall has 2GB of RAM, and has 512MB for /var and 512MB for /tmp, then only 1GB of RAM will be available to the OS for general use.

Special care must be taken when choosing a RAM disk size, which is discussed in the following section.

RAM Disk Sizes

Setting a size too small for /tmp and /var can backfire, especially when it comes to packages. The suggested sizes on the page are an absolute minimum and often much larger sizes are required. The most common failure is that when a package is installed, and parts of the package touch places in both /tmp and /var and it can ultimately fill up the RAM disk and cause other data to be lost. Another common failure is setting /var as a RAM disk and then forgetting to move a squid cache to a location outside of /var - if left unchecked, it will fill up the RAM disk.

For /tmp, a minimum of 40 MiB is required. For /var a minimum of 60 MiB is required. To determine the proper size, check the current usage of the /tmp and /var directories before making a switch. Check the usage several times over the course of a few days so it is not caught at a low point. Watching the usage during a package installation adds another useful data point.

Hard Disk Standby

The Hard disk standby time option activates power management for disk drives in the firewall. The drop-down field sets the number of minutes that the disk can be idle before going into standby mode.

Using standby mode is not necessary for SSD or flash media. For traditional spinning platter hard disks, it may result in power savings and can potentially lengthen the disk lifetime by saving wear, at a cost of slower disk access when resuming from an idle state. Actual results entirely depend on the hardware involved.

The default behavior is *Always On* which prevents the disk from entering standby mode.

Installation Feedback

When this option is set, the firewall will not send its AZTCO-FW Device ID when making requests to AZTCO-FW servers.

3.6.5 System Tunables Tab

The System Tunables tab under System > Advanced provides a means to set run-time FreeBSD system tunables, also known as sysctl OIDs.

Tip: In most cases, the best practice is to leave these tunables at their default values.

Firewall administrators familiar with FreeBSD, or users doing so under the direction of a developer or support representative, may want to adjust or add values on this page so that they will be set as the system starts.

Creating and Editing Tunables

To edit an existing tunable, click ⧉.

To create a new tunable, click ✚ New at the top of the list.

When editing or creating a tunable, the following fields are available:

Tunable The sysctl OID to set

Value The value to which the Tunable will be set.

> Note: Some values have formatting requirements. Due to the vast number of sysctl OIDs, the GUI does not validate that the given Value will work for the chosen Tunable.

Description An optional description for reference.

Click Save when the form is complete.

Tunable OIDs and Values

There are many OIDs available from sysctl, some of them can be set, some are read only outputs, and others must be set before the system boots as Loader Tunables. The full list of OIDs and their possible values is outside the scope of this documentation, but for those interested in digging a little deeper, the sysctl manual page from FreeBSD contains detailed instructions and information.

3.6.6 Notifications

The firewall can notify administrators of important events and errors by displaying an alert in the menu bar, indicated by the 🔔 icon.

In addition to GUI notifications, the firewall also supports the following remote notification methods:

- E-mail using SMTP
- Telegram notification API
- Pushover notification API

General Settings

Certificate Expiration When set, the firewall will issue notifications when certificates approach their expiration date, so that administrators can take corrective action to renew or replace them. Notifications are also sent for expired certificates.

The expiration times are checked daily, and notifications are displayed in the GUI and sent remotely.

Certificate Expiration Threshold The value, in days, at which certificates are considered to be approaching their expiration date.

The default value is currently 27 days. Certificates from Let's Encrypt (*ACME package*) typically renew when they have around 30 days before they expire. The default value is long enough that it does not notify unnecessarily, but with enough time left that problems can be corrected.

Tip: If certificates are imported into the firewall from third party sources which take longer to process, increase this value sufficiently to give administrators enough notice to obtain an updated replacement certificate before the expiration date.

SMTP E-mail

E-mail notifications are delivered by a direct SMTP connection to a mail server. The server must be configured to allow relaying from the firewall or accept authenticated SMTP connections.

Disable SMTP When checked, the firewall will not send SMTP notifications. This is useful to silence notifications while keeping SMTP settings in place for use by other purposes such as packages that utilize e-mail.

E-mail server the hostname or IP address of the e-mail server through which the firewall will send notifications.

SMTP Port of E-mail server the port to use when communicating with the SMTP server. The most common ports are 25 and 587.

In many cases, 25 will not work unless it is to a local or internal mail server. Providers frequently block outbound connections to port 25, so use 587 (the Submission port) when possible.

Connection Timeout to E-Mail Server The length of time, in seconds, that the firewall will wait for an SMTP connection to complete.

Secure SMTP Connection When set, the firewall will attempt an SSL/TLS connection when sending e-mail. The server must accept SSL/TLS connections or support STARTTLS.

Validate SSL/TLS When set, the certificate presented by the mail server is checked for validity against the root certificates trusted by the firewall. Ensuring this validity is the best practice.

In some rare cases a mail server may have a self-signed certificate or a certificate that otherwise fails validation. Unchecking this option will allow notifications to be sent to these servers using SSL/TLS. In this case, communication is still encrypted, but the identity of the server cannot be validated.

From e-mail address the e-mail address for the from: header in notification messages, which specifies the source. Some SMTP servers attempt to validate this address so the best practice is to use a real address in this field. This is commonly set to the same address as Notification E-mail address.

Notification E-mail address the e-mail address for the to: header of the message, which is the destination where the notification e-mails will be delivered by the firewall.

Notification E-Mail Auth Username Optional. If the mail server requires a username and password for authentication, enter the username here.

Notification E-Mail Auth Password Optional. If the mail server requires a username and password for authentication, enter the password here and in the confirmation field.

Notification E-mail Auth Mechanism This field specifies the authentication mechanism required by the mail server. The majority of e-mail servers work with *PLAIN* authentication, others such as MS Exchange may require *LOGIN* style authentication.

Click 💾 Save at the bottom of the page to store the settings before proceeding.

Click ✈ Test SMTP Settings to generate a test notification and send it via SMTP using the previously stored settings. Save settings before clicking this button.

Startup/Shutdown Sound

If the firewall hardware has a PC speaker, it will play a sound when startup finishes and again when a shutdown is initiated.

Check Disable the startup/shutdown beep to prevent the firewall from playing these sounds.

Telegram

The notification system supports the Telegram API which can send notifications to desktops and mobile devices, among others.

Note: Using the Telegram API requires a Telegram Bot and its associated API key.

Enable Telegram When set, the firewall will attempt to send remote notifications using the Telegram API and the settings in this section.

API Key Required. The Telegram Bot API key the firewall will use to authenticate with the Telegram API server.

Chat ID The destination for the notifications. This can be a chat ID number for private notifications, or a channel @username for public notifications.

Click Save at the bottom of the page to store the settings before proceeding.

Click Test Telegram Settings to generate a test notification and send it using the Telegram API with the previously stored settings. Save settings before clicking this button.

Pushover

The notification system supports the Pushover API which can send notifications to desktops and mobile devices, among others.

Note: Using the Pushover API requires a Pushover account user key and API key (Pushover Registration).

Enable Pushover When set, the firewall will attempt to send remote notifications using the Pushover API and the settings in this section.

API Key Required. The Pushover API Key (Pushover Registration) the firewall will use to authenticate with the Pushover API server.

User Key Required. The User Key (Pushover Registration) of the Pushover account to which the API Key belongs.

Notification Sound The notification sound that the end user device (Phone, etc) will play when notification messages are sent by the firewall.

See also:

For a list of sounds and audio, see the Pushover API Notification Sounds Documentation.

Message Priority The message priority for firewall notifications.

Note: For more information about the priorities and their meanings, see the Pushover API Priority Documentation.

The following priorities are available:

Normal Default setting. May trigger sound, vibration, and notification display depending on the user settings and client platform.

Lowest No sound or vibration, but increases the notification count on some platforms.

Low No sound or vibration. May trigger a notification display depending on the user settings and client platform.

High Always play sound and vibrate. Bypasses pre-set quiet hours. Notification display is highlighted in red.

Emergency Similar to *High* priority, but the notification is repeated until acknowledged by the user.

Emergency Priority Notification Retry Interval The amount of time, in seconds, the Pushover servers will send the same notification for *Emergency* priority notifications until the notification is acknowledged.

This parameter must have a value of at least 30 seconds between retries. Default is 60 seconds (1 minute).

Emergency Priority Notification Expiration The duration, in seconds, for which *Emergency* priority notifications will be retried until the notification is acknowledged. Notifications will be resent at intervals determined by the value of Emergency Priority Notification Retry Interval.

This parameter must have a maximum value of at most 10800 seconds (3 hours). Default is 300 seconds (5 minutes).

Click Save at the bottom of the page to store the settings before proceeding.

Click Test Pushover Settings to generate a test notification and send it using the Pushover API with the previously stored settings. Save settings before clicking this button.

3.7 Console Menu Basics

Basic configuration and maintenance tasks can be performed from the AZTCO-FW® system console. The console is available using a keyboard and monitor, serial console, or by using SSH. Access methods vary depending on hardware. Below is an example of what the console menu will look like, but it may vary slightly depending on the version and platform:

```
WAN (wan)          -> vmx0           -> v4/DHCP4: 198.51.100.6/24
                                     v6/DHCP6: 2001:db8::20c:29ff:fe78:6e4e/64
LAN (lan)          -> vmx1           -> v4: 10.6.0.1/24
                                     v6/t6: 2001:db8:1:eea0:20c:29ff:fe78:6e58/64

0)   Logout (SSH only)        8) pfTop
1)   Assign Interfaces        9) Filter Logs
2)   Set interface(s) IP address      10) Restart webConfigurator
3)   Reset webConfigurator password   11) PHP shell + AZTCO-FW tools
4)   Reset to factory defaults 12) Update from console
5)   Reboot system 13) Disable Secure Shell (sshd)
6)   Halt system    14) Restore recent configuration
7)   Ping host      15) Restart PHP-FPM
```

Page Contents

3.7.1 1) Assign Interfaces

This option restarts the Interface Assignment task, which is covered in detail in *Assign Interfaces* and *Manually Assigning Interfaces*. This menu option can create VLAN interfaces, reassign existing interfaces, or assign new ones.

3.7.2 2) Set interface(s) IP address

The script to set an interface IP address can set WAN, LAN, or OPT interface IP addresses, but there are also other useful features of this script:

- The firewall prompts to enable or disable DHCP service for an interface, and to set the DHCP IP address range if it is enabled.
- If the firewall GUI is configured for HTTPS, the menu prompts to switch to HTTP. This helps in cases when the SSL configuration is not functioning properly.
- If the anti-lockout rule on LAN has been disabled, the script enables the anti-lockout rule in case the user has been locked out of the GUI.

3.7.3 3) Reset webConfigurator password

This menu option invokes a script to reset the admin account password and status. The password is reset to the default value of aztco.

The script also takes a few other actions to help regain entry to the firewall:

- If the GUI authentication source is set to a remote server such as RADIUS or LDAP, it prompts to return the authentication source to the Local Database.
- If the admin account has been removed, the script re-creates the account.
- If the admin account is disabled, the script re-enables the account.

3.7.4 4) Reset to factory defaults

This menu choice restores the system configuration to factory defaults. It will also attempt to remove any installed packages.

This action is also available in WebGUI at Diagnostics > Factory Defaults.

See *Resetting to Factory Defaults* for more details about how this process works.

3.7.5 5) Reboot system

This menu choice cleanly shuts down the firewall and restarts the operating system. There are several options which control what the firewall will do when rebooting. The choices offered by the reboot option are explained in *Reboot Methods*. See also:

This action is also available in WebGUI at Diagnostics > Reboot, see *Rebooting the Firewall* for details.

3.7.6 6) Halt system

This menu choice cleanly shuts down the firewall and either halts or powers off, depending on hardware support.

Warning: The best practice is to never cut power from a running system. Halting before removing power is always the safest choice.

See also:

This action is also available in WebGUI at Diagnostics > Halt System. See *Halting and Powering Off the Firewall* for additional details.

3.7.7 7) Ping host

This menu option runs a script which attempts to contact a host to confirm if it is reachable by the firewall through a connected network. The script prompts the user for an IP address, and then the script sends that target host three ICMP echo requests.

The script displays output from the test, including the number of packets received, sequence numbers, response times, and packet loss percentage.

The script uses ping when given an IPv4 address or a hostname, and ping6 when given an IPv6 address.

This is only a basic ping test. For more options, see *Ping Host* to run a similar test from the GUI.

3.7.8 8) pfTop

This menu option invokes pftop which displays a real-time view of the firewall states, and the amount of data they have sent and received. It can help pinpoint sessions currently using large amounts of bandwidth, and may also help diagnose other network connection issues.

See also:

See *Viewing States with pfTop* for more information on how to use pfTop.

3.7.8 9) Filter Logs

The Filter Logs menu option displays firewall log entries in real-time, in their raw form. The raw logs contain much more information per line than the log view in the WebGUI (Status > System Logs, Firewall tab), but not all of this information is easy to read.

Tip: For a simplified console view of the firewall logs in real time with low detail, use the following shell command:

```
clog -f /var/log/filter.log | filterparser.php
```

3.7.10 10) Restart webConfigurator

Restarting the webConfigurator will restart the system process that runs the GUI (nginx). In extremely rare cases the process may have stopped, and restarting it will restore access to the GUI.

If the GUI is not responding and this option does not restore access, invoke menu option 16 to Restart PHP-FPM after using this menu option.

3.7.11 11) PHP shell + AZTCO-FW tools

The PHP shell is a powerful utility that executes PHP code in the context of the running system. As with the normal shell, it is also potentially dangerous to use. This is primarily used by developers and experienced users who are intimately familiar with both PHP and the AZTCO-FW software code base.

Playback Scripts

There are several playback scripts for the PHP Shell that automate simple tasks or enable access to the GUI.

These scripts are run from within the PHP shell like so:

```
AZTCO-FW shell: playback scriptname
```

They may also be run from the command line:

```
# pfSsh.php playback scriptname
```

changepassword

This script changes the password for a user, and also prompts to reset the account properties if it is disabled or expired.

disablecarp / enablecarp

These scripts disable and enable CARP high availability functions, and will deactivate CARP type Virtual IP addresses.

This action does not persist across reboots. **disablecarpmaint / enablecarpmaint**

These scripts disable and enable CARP maintenance mode, which leaves CARP active but demotes this unit so the other node can assume control. This maintenance mode will persist across reboots.

disabledhcpd

This script removes all DHCP configuration from the firewall, effectively disabling the DHCP service and completely

removing all of its settings. **disablereferercheck**

This script disables the HTTP_REFERER check mentioned in *Browser HTTP_REFERER enforcement*. This can help gain access to the GUI if a browser session is triggering this protection.

enableallowallwan

This script adds an allow all rule for IPv4 and IPv6 to the WAN interface.

Warning: Be extremely careful with this option, it is meant to be a temporary measure to gain access to services on the WAN interface of the firewall in situations where the LAN is not usable. Once proper access rules are put in place, remove the rules added by this script.

enablesshd

This script enables the SSH daemon, the same as the console menu option or GUI option.

externalconfiglocator

This script will look for a config.xml file on an external device, such as a USB thumb drive, and will move it in place for use by the firewall.

gatewaystatus

This script prints the current gateway status and statistics. It also accepts an optional parameter brief which prints only the gateway name and status, omitting the addresses and statistical data.

generateguicert

This script creates a new self-signed certificate for the firewall and activates it for use in the GUI. This is useful in cases where the previous certificate is invalid or otherwise not usable. It also fills in the certificate details using the firewall hostname and other custom information, to better identify the host.

gitsync

This complex script synchronizes the PHP and other script sources with files from the AZTCO-FW github repository. It is most useful on development snapshots to pick up changes from more recent commits.

> Warning: This script can be dangerous to use in other circumstances. Only use this under the direction of a knowledgeable developer or support representative.

If the script is run without any parameters it will print a help message outlining its use. More information can be found at *Using gitsync to Update AZTCO-FW Between Snapshots*.

installpkg / listpkg / uninstallpkg

These scripts interface with the package system in a similar way to the GUI. These are primarily used for debugging package issues, comparing information in config.xml compared to the package database.

pfanchordrill

This script recursively searches through pf anchors and prints any NAT or firewall rules it finds. This can help track down unexpected behavior in areas such as UPnP which rely on rules in anchors that are not otherwise visible in the GUI.

pftabledrill

This script prints the contents of all pf tables, which contain addresses used in firewall aliases as well as built-in system tables for features such as bogon network blocking, snort, and GUI/SSH lockout. This script is useful for checking if a specific IP address is found in any table, rather than searching individually.

removepkgconfig

This script removes all traces of package configuration data from the running config.xml. This can be useful if a package has corrupted settings or has otherwise left the packages in an inconsistent state.

removeshaper

This script removes ALTQ traffic shaper settings, which can be useful if the shaper configuration is preventing rules from loading or is otherwise incorrect and preventing proper operation of the firewall.

resetwebgui

This script resets the GUI settings for widgets, dashboard columns, the theme, and other GUI-related settings. It can return the GUI, particularly the dashboard, to a stable state if it is not functioning properly.

restartallwan

This script disables and re-enables each WAN-type interface, which reapplies the interface configuration.

restartdhcpd

This script stops and restarts the DHCP daemon.

restartipsec

This script rewrites and reloads the IPsec configuration for strongSwan.

svc

This script controls the services running on the firewall, similar to interacting with services at Status > Services.

The general form of the command is:

```
playback svc <action> <service name> [service-specific options]
```

The action can be stop, start, or restart.

The service name is the name of the services as found under Status > Services. If the name includes a space, enclose the name in quotes.

The service-specific options vary depending on the service, they are used to uniquely identify services with multiple instances, such as OpenVPN or Captive Portal entries.

Examples:

- Stop miniupnpd:

```
pfSsh.php playback svc stop miniupnpd
```

- Restart OpenVPN client with ID 2:

```
pfSsh.php playback svc restart openvpn client 2
```

- Start the Captive Portal process for zone “MyZone”:

```
pfSsh.php playback svc start captiveportal MyZone
```

3.7.12 12) Upgrade from console

This menu option runs the AZTCO-FW-upgrade script to upgrade the firewall to the latest available version. This is operationally identical to running an upgrade from the GUI and requires a working network connection to reach the update server.

This method of upgrading is covered with more detail in *Upgrading using the Console*.

3.7.13 13) Enable/Disable Secure Shell (sshd)

This option toggles the status of the Secure Shell Daemon, sshd. This option works the same as the option in the WebGUI to enable or disable SSH.

3.7.14 14) Restore recent configuration

This menu option starts a script that lists and restores backups from the configuration history. This is similar to accessing the configuration history from the GUI at Diagnostics > Backup/Restore on the Config History tab (*Restoring from the Config History*).

This script can display the last few configuration files, along with a timestamp and description of the change made in the configuration, the user and IP address that made the change, and the config revision. This is especially useful if a recent configuration error accidentally prevented access to the GUI.

3.7.15 15) Restart PHP-FPM

This menu option stops and restarts the daemon which handles PHP processes for nginx. If the GUI web server process is running but unable to execute PHP scripts, invoke this option. Run this option in conjunction with Restart webConfigurator for the best result.

3.8 Resetting to Factory Defaults

The firewall configuration can be reset back to defaults, a process which also attempts to remove any installed packages. This reset can be performed in the GUI from Diagnostics > Factory Defaults, by using the console menu,

In each case, the firewall will automatically reboot with a default configuration after the reset, which may require console access to resolve.

Note: This process does not remove any changes made to the file system, it only resets the configuration.

If system files have been corrupted or altered in an undesirable way, the best practice is to make a backup and reinstall from installation media.

3.8.1 Factory Default from the GUI

To reset the configuration to factory defaults using the GUI:

- Navigate to Diagnostics > Factory Defaults
- Review the items on the page which will be affected by the reset
- Click ↺ Factory Reset
- Click OK to confirm the action and start the reset process

3.8.2 Factory Default from the Console

To reset the configuration to factory defaults using the console:

- Access the console menu locally or via SSH with an admin-level account (admin, root, or another privileged account using sudo).
- Enter the menu option which corresponds with Reset to factory defaults (e.g. 4)
- Press Enter
- Enter the y to confirm the action
- Press Enter to start the reset process

3.9 Connecting to the GUI

To reach the GUI, follow this basic procedure:

- Connect a client computer to the same network as the LAN interface of the firewall. This computer may be directly connected with a network cable or connected to the same switch as the LAN interface of the firewall.

  By default, the LAN IP address of a new installation of AZTCO-FW software is 192.168.1.1:5687 with a /24 mask (255.255.255.0), and there is also a DHCP server running. If a client computer is set to use DHCP, it should obtain an address in the LAN subnet automatically.

- On the client computer, open a web browser such as Firefox, Safari, or Chrome and navigate to https://192.168.1.1:5687.
  The GUI listens on HTTPS by default, but if the browser attempts to connect using HTTP, it will be redirect by the firewall to the HTTPS port instead.
- Enter the default credentials in the login page:

  username admin

  password aztco

In some cases additional steps may be necessary before the client computer can reach the GUI.

Warning: If the default LAN subnet conflicts with the WAN subnet, the LAN subnet must be changed before connecting it to the rest of the network. Attempting to access the GUI in this situation is unpredictable and unlikely to work until the conflict is resolved.

8.10. Connecting to the GUI

The LAN IP address may be changed and DHCP may be disabled using the console:

- Open the console (VGA, serial, or using SSH from another interface)
- Choose option 2 from the console menu
- Enter the new LAN IP address, subnet mask, and specify whether or not to enable DHCP.
- Enter the starting and ending address of the DHCP pool if DHCP is enabled. This can be any range inside the given subnet.

Note: When assigning a new LAN IP address, it cannot be in the same subnet as the WAN or any other active interface. If there are other devices already present on the LAN subnet, it also cannot be set to the same IP address as an existing host.

If the DHCP server on the firewall is disabled, client computers on LAN must have a statically configured IP address in the LAN subnet, such as 192.168.1.5, with a subnet mask that matches the one given to the firewall, such as 255.255.255.0.

CHAPTER

FOUR

BACKUP AND RECOVERY

4.1 Making Backups in the GUI

Making a backup in the GUI is simple:

- Navigate to Diagnostics > Backup & Restore
- Set any desired options, or leave the options at their default values.
- Click Download Configuration as XML (Figure *GUI Backup*).

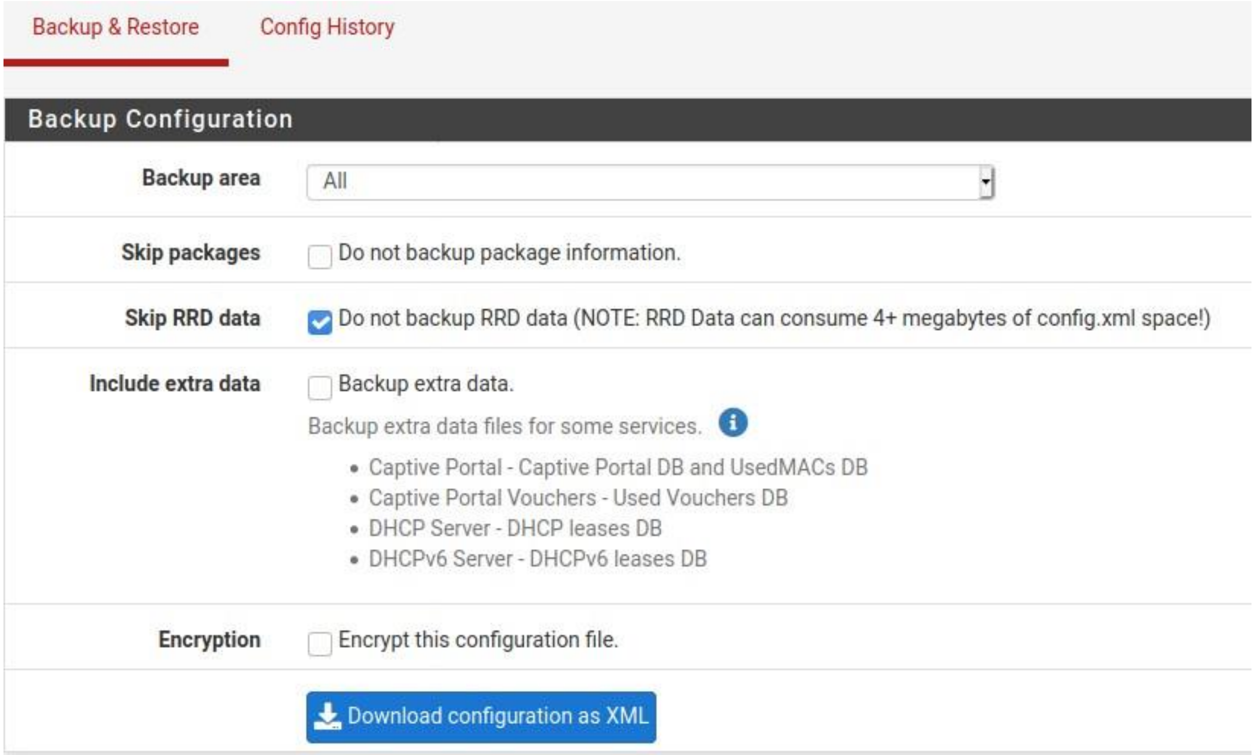

Backup Configuration

Backup area	All
Skip packages	□ Do not backup package information.
Skip RRD data	☑ Do not backup RRD data (NOTE: RRD Data can consume 4+ megabytes of config.xml space!)
Include extra data	□ Backup extra data. Backup extra data files for some services. • Captive Portal - Captive Portal DB and UsedMACs DB • Captive Portal Vouchers - Used Vouchers DB • DHCP Server - DHCP leases DB • DHCPv6 Server - DHCPv6 leases DB
Encryption	□ Encrypt this configuration file.

Fig. 1: GUI Backup

The web browser will then prompt to save the file somewhere on the PC being used to view the GUI. It will be named config-<hostname>-<timestamp>.xml, but that may be changed before saving the file.

4.1.1 Backup Options

When performing a backup, GUI options are available to control what is contained within the backup file.

Backup Area Limits the backup contents to a single configuration area, rather than a complete configuration backup.

> Note: When restoring a configuration containing only a single area, the Restore area value must be set to match.

Skip Packages When set, omits installation data and settings for packages from the backup. This is useful to quickly remove all traces of packages from a configuration.

Skip RRD Data When set (default), the data used to generate monitoring graphs (*Monitoring Graphs*) is exported and included in the backup, so that when the configuration is restored later, the graph data is also restored.

Include Extra Data When set, additional data is stored in the backup file. This includes Captive Portal databases and DHCP lease databases. These databases are volatile, and thus can be useful for transferring to new hosts or for frequent backups, but are not as useful for long-term backups.

Encryption When set, the GUI presents Password and confirmation fields, the contents of which are used by AZTCO-FW® software to encrypt the contents of the backup file with AES-256 before download.

4.2 Restoring from Backups

Backups are not useful without a means to restore them, and by extension, test them. Several means for restoring configurations are available in AZTCO-FW® software. Each method has the same end result: a running firewall identical to when the backup was made.

4.2.1 Restoring with the GUI

The easiest way for most users to restore a configuration is by using the GUI:

- Navigate to Diagnostics > Backup & Restore
- Locate the Restore Backup section (Figure *GUI Restore*).
- Select the area to restore, or leave at the default selection for a complete backup.

> Note: This value must match the Backup area chosen when creating the backup.

- Click Browse
- Locate the backup file on the local PC
- Click Restore Configuration

The firewall will then apply the configuration and reboot with the settings obtained from the backup file.

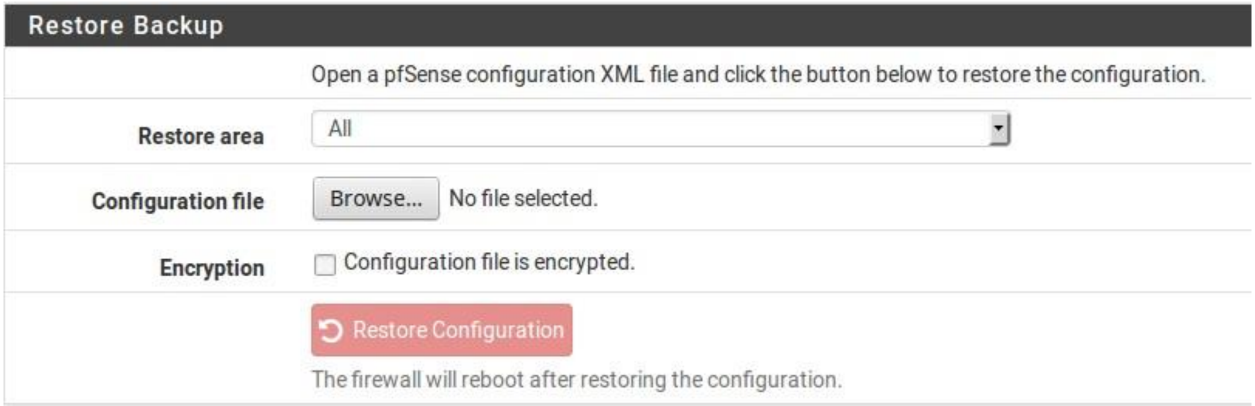

Restore Backup

Open a pfSense configuration XML file and click the button below to restore the configuration.

Restore area: All

Configuration file: Browse... No file selected.

Encryption: □ Configuration file is encrypted.

Restore Configuration

The firewall will reboot after restoring the configuration.

Fig. 2: GUI Restore

While easy to work with, this method has prerequisites when dealing with a full restore to a new installation. First, it would need to be done after the new target system is fully installed and running. Second, it requires an additional PC connected to a working network or crossover cable behind the firewall being restored.

Restore Options

Restore Area Restores a backup containing only a single configuration area, rather than a complete configuration backup.

> Warning: This does not restore one area from a full backup, the backup file must only contain the area to restore.

> Note: This value must match the Backup area chosen when creating the backup.

Configuration File A Browse button to select a backup file to upload and restore.

Preserve Switch Configuration This option is available on AZTCO-FW hardware with integrated switches. When set, the current active switch configuration will be copied into the restored configuration, preserving it for later use. This makes it easier to restore a configuration from hardware without an integrated switch.

> Note: This only copies the integrated switch configuration, and does not copy VLAN or LAGG interface entries which may be relevant to using the switch. This behavior is safer, as the configuration being restored may also contain important configuration data in those areas.

Encryption When set, a Password field is presented, the contents of which is used by the firewall to decrypt the contents of the backup file before restoring the configuration.

4.2.2 Restoring from the Config History

For minor problems, using one of the internal backups on the firewall is the easiest way to back out a change. The previous 30 configurations are stored in the Configuration History, along with the current running configuration.

Each row in the configuration history list shows the date the configuration file was made, the configuration version, the user and IP address of a person making a change in the GUI, the page that made the change, and in some cases, a brief description of the change that was made. The action buttons to the right of each row show a description of what they do when the mouse pointer is hovered over the button.

To restore a configuration from the history:

- Navigate to Diagnostics > Backup & Restore
- Click the Config History tab (Figure *Configuration History*)
- Locate the desired backup in the list
- Click ⟲ to restore that configuration file

Restoring a configuration with this method does not initiate an automatic reboot. Minor changes do not require a reboot, though reverting some major changes will.

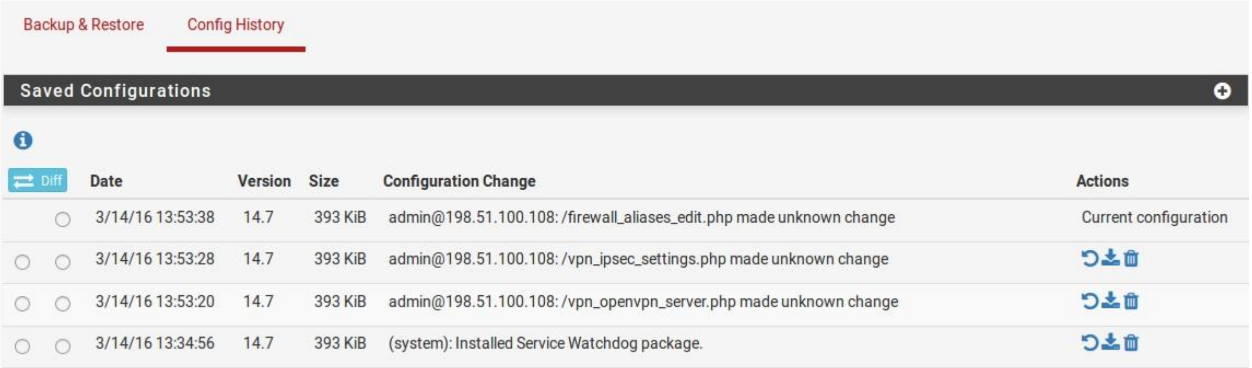

⇄ Diff		Date	Version	Size	Configuration Change	Actions
	○	3/14/16 13:53:38	14.7	393 KiB	admin@198.51.100.108: /firewall_aliases_edit.php made unknown change	Current configuration
○	○	3/14/16 13:53:28	14.7	393 KiB	admin@198.51.100.108: /vpn_ipsec_settings.php made unknown change	⟲⬇🗑
○	○	3/14/16 13:53:20	14.7	393 KiB	admin@198.51.100.108: /vpn_openvpn_server.php made unknown change	⟲⬇🗑
○	○	3/14/16 13:34:56	14.7	393 KiB	(system): Installed Service Watchdog package.	⟲⬇🗑

Fig. 3: Configuration History

If a change was only made in one specific section, such as firewall rules, trigger a refresh in that area of the GUI to enable the changes. For firewall rules, a filter reload would be sufficient. For OpenVPN, edit and save the VPN instance. The necessary actions to take depend on the changes in the restored configuration, but the best way ensure that the full configuration is active is to reboot.

If necessary, reboot the firewall with the new configuration by going to Diagnostics > Reboot System and click Yes.

Previously saved configurations may be deleted by clicking 🗑, but do not delete them by hand to save space; the old configuration backups are automatically deleted when new ones are created. It is desirable to remove a backup from a known-bad configuration change to ensure that it is not accidentally restored.

A copy of the previous configuration may be downloaded by clicking ⬇.

Configuration Backup Cache Settings

The amount of backups stored in the configuration history may be changed if needed.

- Navigate to Diagnostics > Backup & Restore
- Click the Config History tab
- Click ⊕ at the right end of the Configuration Backup Cache Settings bar to expand the settings
- Enter the new number of configurations to retain in the Backup Count field
- Click Save

Along with the configuration count, the page also displays the amount of space consumed by the backup cache.

Config History Diff

The differences between any two configuration files may be viewed in the Config History tab. To the left of the configuration file list there are two columns of radio buttons. Use the leftmost column to select the *older* of the two configuration files, and then use the right column to select the *newer* of the two files. Once both files have been selected, click Diff at either the top or bottom of the column.

Console Configuration History

The configuration history is also available from the console menu as option 15, Restore Recent Configuration. The menu selection will list recent configuration files and offer to restore one. This is useful if a recent change has locked administrators out of the GUI or taken the firewall off the network.

4.2.3 Restoring by Mounting the Disk

Attaching the disk from an installation of AZTCO-FW software to a computer running FreeBSD enables the drive to be mounted by the FreeBSD host and a new configuration may be copied directly onto the installed system, or a configuration file from a failed system may be copied off.

Note: This can also be performed on a separate installation of AZTCO-FW in place of a computer running FreeBSD, but do not use an active production firewall for this purpose. Instead, use a spare or test firewall.

The config.xml file is kept in /cf/conf/, but the difference is in the location where this directory resides. This is part of the root slice (typically da0p2). The drive and partition name will vary depending on disk type and position in the host.

4.3 Restoring a Configuration File to a Different Version

Configurations are specific to a given version of AZTCO-FW® software. The configuration is the same on all platforms and architectures using the same version of AZTCO-FW software. The version of FreeBSD used is not relevant.

Generally speaking, a complete older configuration version can always be restored to a newer release of AZTCO-FW software. The firewall will upgrade the configuration as needed provided that has the entire configuration and not a partial copy.

A newer configuration cannot be restored to an older release that had a different configuration version. Certain releases of AZTCO-FW software had the same configuration version, and restoring between those is possible, but still not recommended. See *Versions of AZTCO-FW and FreeBSD* to see which configuration versions were used on specific releases.

A configuration section or partial configuration cannot be restored between different configuration versions. It may work by pure luck, but often there are configuration format differences that require changes to be made to the older configuration. These changes are automatic if a complete configuration is restored. If a partial restore is required, perform a full upgrade in a test VM or lab and then copy the needed section out of the resulting config.xml post-upgrade.

4.4 Password Storage Security Policies

Sensitive data such as PPPoE/PPTP client, PPTP VPN, DynDNS passwords as well as remote authentication servers RADIUS (shared secret), LDAP (bind user password), and IPsec shared secrets, among others, appear in plain text or with reversible Base64 encoding in the AZTCO-FW® software configuration file, config.xml. This is a deliberate design decision in m0n0wall that has been carried over here.

Since the firewall cannot prompt the user for a password each time it is required, the implementations of affected areas require plain text passwords to operate. AZTCO-FW software could, of course, use some snake oil encryption on those passwords, but that would only create a false sense of security. Any encryption applied to the passwords could be reversed by anyone with access to the source code (i.e. everybody). Hashes like SHA256 cannot be used where the plain text password is needed at a later stage, unlike for the system password, which is only stored as a hash.

By leaving the passwords in plain text, it is very clear that config.xml deserves to be stored in a secure location (and/or encrypted with one of the countless programs out there). Any sort of hashing used would not be secure, and would be dangerous because it would give the impression of security where none exists.

See also:

- *Backup Files and Directories with the Backup Package*

Thanks to the XML-based configuration file used by AZTCO-FW® software, backups are a breeze. All of the settings for the system are held in one single file (see *XML Configuration File*). In the vast majority of cases, this one file can be used to restore a system to a fully working state identical to what was running previously. There is no need to make an entire system backup, as the base system files are not modified by a normal, running, system.

Note: In rare cases, packages may store files outside of config.xml, check the package documentation for additional information and backup suggestions.

FIVE

INTERFACE TYPES AND CONFIGURATION

5.1 Interface Configuration

To assign a new interface:

- Navigate to Interfaces > Assignments
- Pick the new interface from the Available network ports list
- Click + Add

The newly assign interface will be shown in the list. The new interface will have a default name allocated by the firewall such as OPT1 or OPT2, with the number increasing based on its assignment order. The first two interfaces default to the names WAN and LAN but they can be renamed. These OPTx names appear under the Interfaces menu, such as Interfaces > OPT1. Selecting the menu option for the interface will open the configuration page for that interface.

The following options are available for all interface types.

5.1.1 Description

The name of the interface. This will change the name of the interface on the Interfaces menu, on the tabs under Firewall > Rules, under Services > DHCP, and elsewhere throughout the GUI. Interface names may only contain letters, numbers and the only special character that is allowed is an underscore (“_”). Using a custom name makes it easier to remember the purpose of an interface and to identify an interface for adding firewall rules or choosing other per-interface functionality.

10.1.2 IPv4 Configuration Type

Configures the IPv4 settings for the interface. Details for this option are in the next section, *IPv4 WAN Types*.

10.1.3 IPv6 Configuration Type

Configures the IPv6 settings for the interface. Details for this option are in *IPv6 WAN Types*.

5.1.2 MAC address

The MAC address of an interface can be changed (“spoofed”) to mimic a previous piece of equipment.

Warning: We recommend avoiding this practice. The old MAC would generally be cleared out by resetting the equipment to which this firewall connects, or by clearing the ARP table, or waiting for the old ARP entries to expire. It is a long-term solution to a temporary problem.

Spoofing the MAC address of the previous firewall can allow for a smooth transition from an old router to a new router, so that ARP caches on devices and upstream routers are not a concern. It can also be used to fool a piece of equipment into believing that it's talking to the same device that it was talking to before, as in cases where a certain network router is using static ARP or otherwise filters based on MAC address. This is common on cable modems, where they may require the MAC address to be registered if it changes.

One downside to spoofing the MAC address is that unless the old piece of equipment is permanently retired, there is a risk of later having a MAC address conflict on the network, which can lead to connectivity problems. ARP cache problems tend to be very temporary, resolving automatically within minutes or by power cycling other equipment.

If the old MAC address must be restored, this option must be emptied out and then the firewall *must* be rebooted. Alternately, enter the original MAC address of the network card and save/apply, then empty the value again.

5.1.3 MTU (Maximum Transmission Unit)

The Maximum Transmission Unit (MTU) size field can typically be left blank, but can be changed when required. Some situations may call for a lower MTU to ensure packets are sized appropriately for an Internet connection. In most cases, the default assumed values for the WAN connection type will work properly. It can be increased for those using jumbo frames on their network.

On a typical Ethernet style network, the default value is 1500, but the actual value can vary depending on the interface configuration.

5.1.4 MSS (Maximum Segment Size)

Similar to the MTU field, the MSS field "clamps" the Maximum Segment Size (MSS) of TCP connections to the specified size in order to work around issues with Path MTU Discovery.

5.1.5 Speed and Duplex

The default value for link speed and duplex is to let the firewall decide what is best. That option typically defaults to *Autoselect*, which negotiates the best possible speed and duplex settings with the peer, typically a switch.

The speed and duplex setting on an interface must match the device to which it is connected. For example, when the firewall is set to *Autoselect*, the switch must also be configured for *Autoselect*. If the switch or other device has a specific speed and duplex forced, it must be matched by the firewall.

5.1.6 Block Private Networks

When Block private networks is active, AZTCO-FW® software inserts a rule automatically that prevents any RFC 1918 networks (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16) and loopback (127.0.0.0/8) from communicating on that interface. This option is usually only desirable on WAN type interfaces to prevent the possibility of privately numbered traffic coming in over a public interface.

5.1.7 Block bogon networks

When Block bogon networks is active, AZTCO-FW software will block traffic from a list of unallocated and reserved networks. This list is periodically updated by the firewall automatically.

> Warning: This option should only be used on external interfaces (WANs), it is not necessary on local interfaces and it can potentially block required local traffic.

See *Block Bogon Networks* for more details on how this feature works.

5.2 IPv4 WAN Types

Once an interface has been assigned, in most cases it will require an IP address. For IPv4 connections, the following choices are available: Static IPv4, DHCP, PPP, PPPoE, PPTP, and L2TP. These options are selected using the IPv4 Configuration Type selector on an interface page (e.g. Interfaces > WAN).

5.2.1 None

When IPv4 Configuration Type is set to *None*, IPv4 is disabled on the interface. This is useful if the interface has no IPv4 connectivity or if the IP address on the interface is being managed in some other way, such as for an OpenVPN or GIF interface.

5.2.2 Static IPv4

With Static IPv4, the interface contains a manually configured IP address. When chosen, three additional fields are available on the interface configuration screen: IPv4 Address, a CIDR subnet mask selector, and the IPv4 Upstream Gateway field.

To configure the interface for static IPv4 on an internal interface (e.g. LAN, DMZ):

- Select *Static IPv4* under IPv4 Configuration Type
- Enter the IPv4 address for the interface into the IPv4 address box
- Choose the appropriate subnet mask from the CIDR drop-down after the address box
- Do not select an IPv4 Upstream Gateway

To configure the interface for static IPv4 on a WAN type interface:

- Select *Static IPv4* under IPv4 Configuration Type

- Enter the IPv4 address for the interface into the IPv4 address box
- Choose the appropriate subnet mask from the CIDR drop-down after the address box

10.2. IPv4 WAN Types

- Perform one of the following to use a gateway on the interface:
  - Select an IPv4 Upstream Gateway from the list, *OR*
  - Click + Add a new gateway to create a new gateway if one does not already exist. Clicking that button displays a modal form to add the gateway without leaving this page. Fill in the details requested on the new form:

    Default Gateway If this is the only WAN or will be a new default WAN, check this box. The default IPv4 and IPv6 gateways work independently of one another. The two need not be on the same circuit. Changing the default IPv4 gateway has no effect on the IPv6 gateway, and vice versa.

    Gateway Name The name used to refer to the gateway internally, as well as in places like Gateway Groups, the Quality Graphs, and elsewhere.

    Gateway IPv4 The IP address of the gateway. This address must be inside of the same subnet as the Static IPv4 address when using this form.

    Description A bit of text to indicate the purpose of the gateway.

    * Click + Add

Note: Selecting an IPv4 Upstream Gateway from the drop-down list or adding and selecting a new gateway will make AZTCO-FW® treat this interface as a WAN type interface for NAT and related functions. This is not desirable for internal interfaces such as LAN or a DMZ. Gateways may still be used on internal interfaces for the purpose of static routes without selecting an IPv4 Upstream Gateway here on the interfaces screen.

5.2.3 DHCP

When an interface is set to DHCP, AZTCO-FW will attempt automatic IPv4 configuration of this interface via DHCP. This option also activates several additional fields on the page. Under most circumstances these additional fields may be left blank.

Hostname Some ISPs require the Hostname for client identification. The value in the Hostname field is sent as the DHCP client identifier and hostname when requesting a DHCP lease.

Alias IPv4 Address This value used as a fixed IPv4 alias address by the DHCP client since a typical IP Alias VIP cannot be used with DHCP. This can be useful for accessing a piece of gear on a separate, statically numbered network outside of the DHCP scope. One example would be for reaching a cable modem management IP address.

Reject Leases From An IPv4 address for a DHCP server that should be ignored. For example, a cable modem that hands out private IP addresses when the cable sync has been lost. Enter the private IP address of the modem here, e.g. 192.168.100.1 and the firewall will never pick up or attempt to use a an IP address supplied by the specified server.

Advanced Configuration Enables options to control the protocol timing. In the vast majority of cases this must be left unchecked and the options inside unchanged.

Protocol Timing The fields in this area give fine-grained control over the timing used by dhclient when managing an address on this interface. These options are almost always left at their default values. For more details on what each field controls, see the dhclient man page

Presets Has several options for preset protocol timing values. These are useful as a starting point for custom adjustments or for use when the values need to be reset back to default values.

Configuration Override Enables a field to use a custom dhclient configuration file. The full path must be given. Using a custom file is rarely needed, but some ISPs require DHCP fields or options that are not supported in the AZTCO-FW GUI.

5.2.4 PPP Types

The various PPP-based connection types such as PPP, PPPoE, PPTP, and L2TP are all covered in detail earlier in this chapter (*PPPs*). When one of these types is selected here on the interfaces screen, their basic options can be changed as described. To access the advanced options, follow the link on this page or navigate to Interfaces > Assignments on the PPPs tab, find the entry, and edit it there.

5.3 IPv6 WAN Types

Similar to IPv4, the IPv6 Configuration Type controls if and how an IPv6 address is assigned to an interface. There are several different ways to configure IPv6 and the exact method depends on the network to which this firewall is connected and how the ISP has deployed IPv6.

Warning: Every ISP is different and large providers can even vary by region.

The ISP determines IPv6 settings for a circuit, and they are the only valid source for that information. As such, this documentation does not include examples for specific providers. Contact the ISP for information about their IPv6 client settings and requirements.

The ISP should provide instructions and specific values for configuring IPv6 on their service. For example, on a circuit with a static IPv6 configuration the ISP should supply the subnet addresses and prefix values for the WAN itself, as well as for routed prefixes. Providers who require DHCPv6 should supply values for settings such as the prefix delegation size, along with any requirements they have for client behavior.

See also:

For more information on IPv6, including a basic introduction, see *IPv6*.

5.3.1 None

When IPv6 Configuration Type is set to *None*, IPv6 is disabled on the interface. This is useful if the interface has no IPv6 connectivity or if the IP address on the interface is being managed in some other way, such as for an OpenVPN or GIF interface.

5.3.2 Static IPv6

The Static IPv6 controls work identically to the Static IPv4 settings. See *Static IPv4* for details.

With Static IPv6, the interface contains a manually configured IPv6 address. When chosen, three additional fields are available on the interface configuration screen: IPv6 Address, a prefix length selector, and the IPv6 Upstream Gateway field.

The default IPv4 and IPv6 gateways work independently of one another. The two need not be on the same circuit. Changing the default IPv4 gateway has no effect on the IPv6 gateway, and vice versa.

5.3.3 DHCP6

DHCP6 configures AZTCO-FW® software to attempt automatic IPv6 configuration of this interface via DHCPv6. DHCPv6 will configure the interface with an IP address, prefix length, DNS servers, etc. but not a gateway. The gateway is obtained via router advertisements, so this interface will be set to accept router advertisements. This is a design choice as part of the IPv6 specification, not a limitation of AZTCO-FW. For more information on router advertisements, see *Router Advertisements*.

Several additional fields are available for IPv6 DHCP that do not exist for IPv4 DHCP:

Use IPv4 Connectivity as Parent Interface When set, the IPv6 DHCP request is sent using IPv4 on this interface, rather than using native IPv6. This is only required in special cases when the ISP requires this type of configuration.

Request only an IPv6 Prefix When set, the DHCPv6 client does not request an address for the interface itself, it only requests a delegated prefix.

DHCPv6 Prefix Delegation Size If the ISP supplies a routed IPv6 network via prefix delegation, they will publish the delegation size, which can be selected here. It is typically a value somewhere between *48* and *64*. For more information on how DHCPv6 prefix delegation works, see *DHCP6 Prefix Delegation*. To use this delegation, another internal interface must be set to an IPv6 Configuration Type of *Track Interface* (*Track Interface*) so that it can use the addresses delegated by the upstream DHCPv6 server.

Send IPv6 Prefix Hint When set, the DHCPv6 Prefix Delegation Size is sent along with the request to inform the upstream server how large of a delegation is desired by this firewall. If an ISP allows the choice, and the chosen size is within their allowed range, the requested size will be given instead of the default size.

Debug When set, the DHCPv6 client is started in debug mode.

Advanced Configuration Enables a wide array of advanced tuning parameters for the DHCPv6 client. These options are rarely used, and when they are required, the values are dictated by the ISP or network administrator. See the dhcp6c.conf man page for details.

Configuration Override Enables a field to use a custom configuration file. The full path must be given. Using a custom file is rarely needed, but some ISPs require DHCP fields or options that are not supported in the AZTCO-FW WebGUI.

5.3.4 SLAAC

Stateless address autoconfiguration (*SLAAC*) as the IPv6 type makes AZTCO-FW attempt to configure the IPv6 address for the interface from router advertisements (RA) that advertise the prefix and related information. Note that DNS is not typically provided via RA, so AZTCO-FW will still attempt to get the DNS servers via DHCPv6 when using SLAAC. In the future, the RDNSS extensions to the RA process may allow DNS servers to be obtained from RA. For more information on router advertisements, see *Router Advertisements*.

5.3.5 6RD Tunnel

6RD is an IPv6 tunneling technology employed by some ISPs to quickly enable IPv6 support for their networks, passing IPv6 traffic inside specially crafted IPv4 packets between and end user router and the ISP relay. It is related to 6to4 but is intended to be used within the ISP network, using the IPv6 addresses from the ISP for client traffic. To use 6RD, the ISP must supply three pieces of information: The 6RD prefix, the 6RD Border Relay, and the 6RD IPv4 Prefix length.

6RD Prefix The 6RD IPv6 prefix assigned by the ISP, such as 2001:db8::/32.

6RD Border Relay The IPv4 address of the ISP 6RD relay.

6RD IPv4 Prefix Length Controls how much of the end user IPv4 address is encoded inside of the 6RD prefix. This is normally supplied by the ISP. A value of 0 means the entire IPv4 address will be embedded inside the 6RD prefix. This value allows ISPs to effectively route more IPv6 addresses to customers by removing redundant IPv4 information if an ISP allocation is entirely within the same larger subnet.

5.3.6 6to4 Tunnel

Similar to 6RD, *6to4* is another method of tunneling IPv6 traffic inside IPv4. Unlike 6RD, however, 6to4 uses constant prefixes and relays. As such there are no user-adjustable settings for using the *6to4* option. The 6to4 prefix is always 2002::/16. Any address inside of the 2002::/16 prefix is considered a 6to4 address rather than a native IPv6 address. Also unlike 6RD, a 6to4 tunnel can be terminated anywhere on the Internet, not only at the end user ISP, so the quality of the connection between the user and the 6to4 relay can vary widely.

6to4 tunnels are always terminated at the IPv4 address of 192.88.99.1. This IPv4 address is anycasted, meaning that although the IPv4 address is the same everywhere, it can be routed regionally toward a node close to the user.

Another deficiency of 6to4 is that it relies upon other routers to relay traffic between the 6to4 network and the remainder of the IPv6 network. There is a possibility that some IPv6 peers may not have connectivity to the 6to4 network, and thus these would be unreachable by clients connecting to 6to4 relays, and this could also vary depending upon the 6to4 node to which the user is actually connected.

5.3.7 Track Interface

The *Track Interface* choice works in concert with another IPv6 interface using DHCPv6 Prefix Delegation. When a delegation is received from the ISP, this option designates which interface will be assigned the IPv6 addresses delegated by the ISP and in cases where a larger delegation is obtained, which prefix inside the delegation is used.

IPv6 Interface A list of all interfaces on the system currently set for dynamic IPv6 WAN types offering prefix delegation (DHCPv6, PPPoE, 6rd, etc.). Select the interface from the list which will receive the delegated subnet information from the ISP.

IPv6 Prefix ID If the ISP has delegated more than one prefix via DHCPv6, the IPv6 Prefix ID controls which of the delegated /64 subnets will be used on this interface. This value is specified in hexadecimal.

For example, If a /60 delegation is supplied by the ISP that means 16 /64 networks are available, so prefix IDs from 0 through f may be used.

For more information on how prefix delegation works, see *DHCP6 Prefix Delegation*.

5.4 Interface Groups

Unlike the other interfaces in this chapter, an Interface Group is not a type of interface that can be assigned. Interface groups are used to apply firewall or NAT rules to a set of interfaces on a common tab. If this concept is unfamiliar, consider how the firewall rules for OpenVPN, the PPPoE server, or L2TP server work. There are multiple interfaces in the underlying OS, but the rules for all of them are managed on a single tab for each type. If many interfaces of a similar function are present on the firewall that need practically identical rules, an interface group may be created to add rules to all of the interfaces at the same time. Interfaces can still have their own individual rules, which are processed after the group rules.

To create an interface group:

- Navigate to Interfaces > Assignments, Interface Groups tab
- Click + Add to create a new group
- Enter a Group Name. This name may only contain upper and lowercase letters, no numbers, spaces, or special characters
- Enter a Group Description (optional)
- Add interfaces as Group Members by ctrl-clicking to select entries from the interface list • Click Save

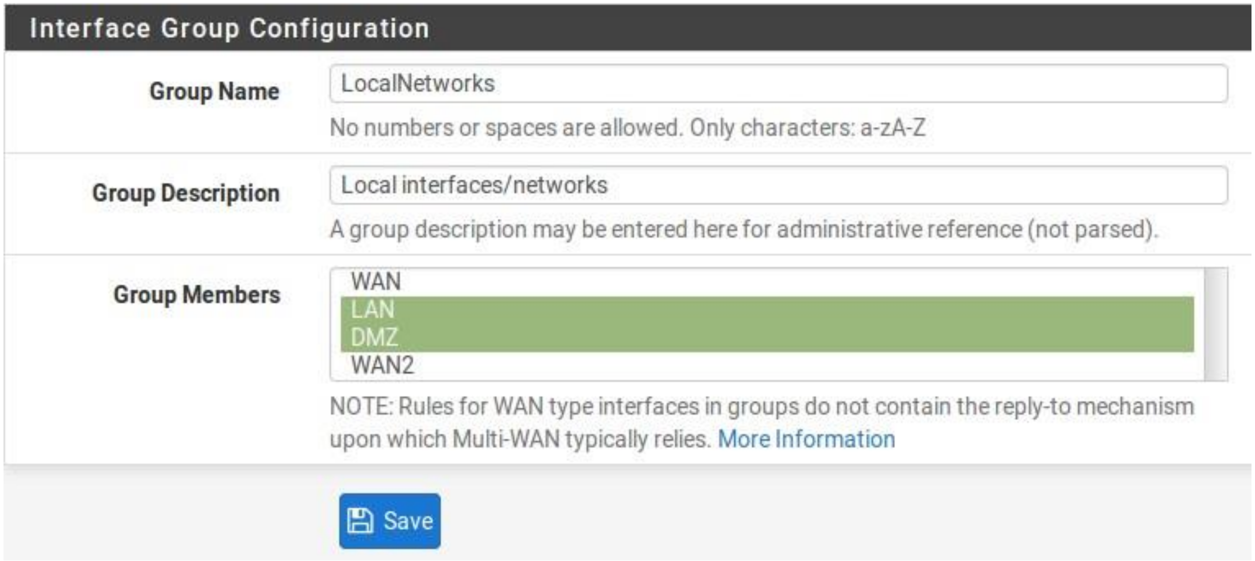

Interface Group Configuration

Group Name: LocalNetworks

No numbers or spaces are allowed. Only characters: a-zA-Z

Group Description: Local interfaces/networks

A group description may be entered here for administrative reference (not parsed).

Group Members: WAN, LAN, DMZ, WAN2

NOTE: Rules for WAN type interfaces in groups do not contain the reply-to mechanism upon which Multi-WAN typically relies. More Information

Fig. 1: Add Interface Group

Interface groups each have an individual tab under Firewall > Rules to manage their rules. Figure *Interface Group Firewall Rules Tab* shows the firewall rule tab for the group defined in figure *Add Interface Group*

See also:

Configuring firewall rules for information on managing firewall rules.

10.4. Interface Groups

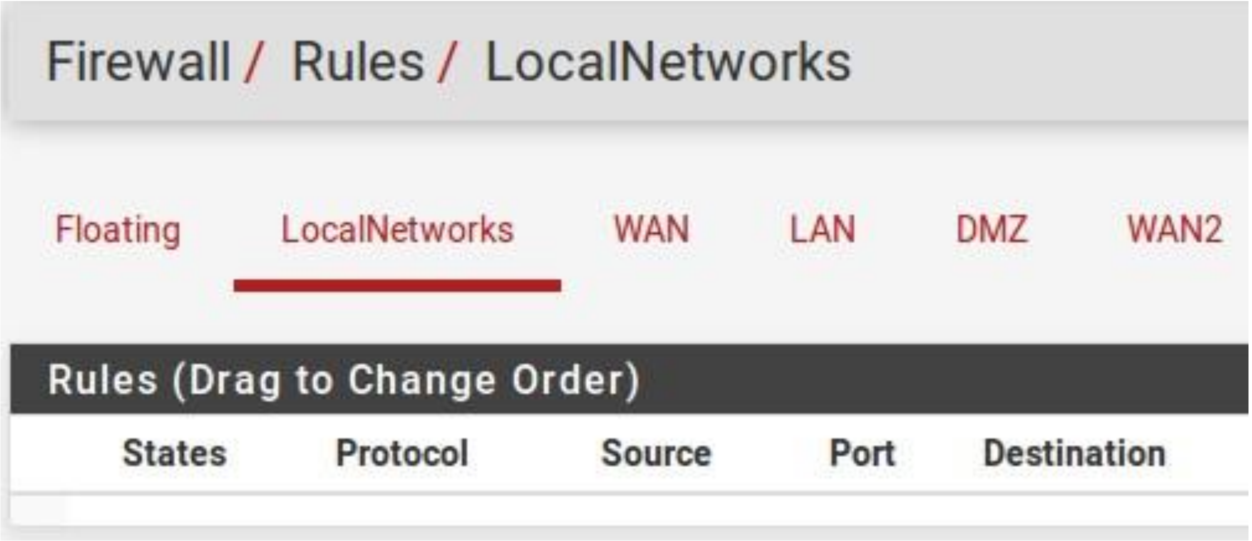

Firewall / Rules / LocalNetworks

Fig. 2: Interface Group Firewall Rules Tab

5.4.1 Group Rule Processing Order

The rule processing order for user rules is:

- Floating rules
- Interface group rules
- Rules on the interface directly

For example, if a rule on the group tab matches a connection, the interface tab rules will not be consulted. Similarly, if a floating rule with Quick set matched a connection, the interface group rules will not be consulted.

The processing order prevents some combination of rules that otherwise might be a good fit. For example, if a general blocking rule is present on the group, it cannot be overriden by a rule on a specific interface. Same with a pass rule, a specific interface rule cannot block traffic passed on a group tab rule.

5.4.2 Use with WAN Interfaces

We do not recommend using interface groups with multiple WANs. Doing so may appear to be convenient, but the group rules do not receive the same treatment as actual WAN tab rules. For example, rules on a tab for a WAN-type interface (Gateway selected on the interface configuration) will receive reply-to which allows pf to return traffic back via the interface from which it entered. Group tab rules do not receive reply-to which effectively means that the group rules only function as expected on the WAN with the default gateway.

10.4. Interface Groups

5.5 PPPs

There are four types of PPP interfaces:

- Plain PPP for 3G/4G and modem devices
- PPPoE for DSL or similar connections
- PPTP and L2TP for ISPs that require them for authentication.

In most cases these are managed from the interface settings directly, but they can also be edited under Interfaces > Assignments on the PPPs tab.

See also:

PPP Logs

5.5.1 Multi-Link PPP (MLPPP)

Editing a PPP instance also allows Multi-Link PPP (MLPPP) to be configured for supported providers. MLPPP bonds multiple PPP links into a single larger aggregate channel. Unlike other multi-WAN techniques, with MLPPP it is possible to use the full bandwidth of all links for a single connection, and the usual concerns about load balancing and failover do not apply. The MLPPP link is presented as one interface with one IP address, and if one link fails, the connection functions the same but with reduced capacity.

For more information on MLPPP, see *Multiple WAN Connections*.

5.5.2 PPP (Point-to-Point Protocol) Interface Types

Add or edit a PPP entry as follows:

- Navigate to Interfaces > Assignments on the PPPs tab
- Click ⧉ to edit an existing entry or ✚ to add a new entry
- Set the Link Type, which changes the remaining options on the page. The link types are explained throughout the remainder of this section.

PPP (3G/4G, Modem)

The PPP link type is used for talking to a modem over a serial device. This can be anything from a USB 3G/4G dongle for accessing a cellular network down to an old hardware modem for dial-up access. Upon selecting the PPP link type, the Link Interface(s) list is populated with serial devices that can be used to communicate with a modem. Click on a specific entry to select it for use. After selecting the interface, optionally enter a Description for the PPP entry.

Note: The serial device for a modem is not automatically detected. Some modems present themselves as several devices, and the subdevice for the PPP line may be any of the available choices, but start with the last device, then try the first, and then others in between if none of those function.

When configuring a 3G/4G network, the Service Provider options pre-fill other relevant fields on the page.

- Select a Country, such as *United States*, to activate the Provider drop-down with known cellular providers in that country
- Select a Provider from the list, such as *T-Mobile*, to activate the Plan drop-down.
- Select a Plan and the remaining fields will be filled with known values for that Provider and Plan

The Service Provider options can be configured manually if other values are needed, or when using a provider that is not listed:

Username and Password The credentials used for the PPP login.

Phone Number The number to dial at the ISP to gain access. For 3G/4G this tends to be a number such as 99# or #777, and for dial-up this is usually a traditional telephone phone number.

Access Point Name (APN) This field is required by some ISPs to identify the service to which the client connects. Some providers use this to distinguish between consumer and business plans, or legacy networks.

APN Number Optional setting. Defaults to 1 if the APN is set, and ignored when APN is unset.

SIM PIN Security code on the SIM to prevent unauthorized use of the card. Do not enter anything here if the SIM does not have a PIN.

SIM PIN Wait Number of seconds to wait for SIM to discover network after the PIN is sent to the SIM. If the delay is not long enough, the SIM may not have time to initialize properly after unlocking.

Init String The modem initialization string, if necessary. Do not include AT at the beginning of the command. Most modern modems do not require a custom initialization string.

Connection Timeout Time to wait for a connection attempt to succeed, in seconds. Default is 45 seconds.

Uptime Logging When checked, the uptime for the connection is tracked and displayed on Status > Interfaces.

PPPoE (Point-to-Point Protocol over Ethernet)

PPPoE is a popular method of authenticating and gaining access to an ISP network, most commonly found on DSL networks.

To configure a PPPoE link, start by setting Link Type to *PPPoE* and complete the remainder of the settings as follows:

Link Interface(s) A list network interfaces that can be used for PPPoE. These are typically physical interfaces but it can also work over some other interface types such as VLANs. Select one for normal PPPoE, or multiple for MLPPP.

Description An optional text description of the PPP entry

Username and Password The credentials for this PPPoE circuit. These will be provided by the ISP, and the username is typically in the form of an e-mail address, such as mycompany@ispexample. com.

Service Name Left blank for most ISPs, some require this to be set to a specific value. Contact the ISP to confirm the value if the connection does not function when left blank.

Configure NULL Service Name Some ISPs require NULL be sent instead of a blank service name. Check this option when the ISP considers this behavior necessary.

Periodic Reset Configures a pre-set time when the connection will be dropped and restarted. This is rarely needed, but in certain cases it can better handle reconnections when an ISP has forced daily reconnections or similar quirky behavior.

Warning: Due to limitations in the way PPPoE frames are processed by network cards, incoming PPPoE traffic is limited to a single network interface queue. As such, performance may be limited or otherwise lower than expected. See *PPPoE with Multi-Queue NICs* for details.

PPTP (Point-to-Point Tunneling Protocol)

Not to be confused with a PPTP VPN, this type of PPTP interface is meant to connect to an ISP and authenticate, much the same as PPPoE works. The options for a PPTP WAN are identical to the PPPoE options of the same name. Refer to the previous section for configuration information.

L2TP (Layer 2 Tunneling Protocol)

L2TP, as it is configured here, is used for connecting to an ISP that requires it for authentication as a type of WAN. L2TP works identically to PPTP. Refer to the previous sections for configuration information.

5.5.3 Advanced PPP Options

All PPP types have several advanced options in common that can be edited in their entries here. In most cases these settings need not be altered. To show these options, click ⚙ Display Advanced.

Dial On Demand The default behavior for a PPP link is to immediately connect and it will immediately attempt to reconnect when a link is lost. This behavior is described as Always On. Dial-on-Demand will delay this connection attempt. When set, the firewall will wait until a packet attempts to leave the via this interface, and then it will connect. Once connected, it will not automatically disconnect.

Idle Timeout A PPP connection will be held open indefinitely by default. A value in Idle Timeout, specified in seconds, will cause the firewall to monitor the line for activity. If there is no traffic on the link for the given amount of time, the link will be disconnected. If Dial-on-Demand has also been set, the firewall will return to dial-on-demand mode.

Note: AZTCO-FW® software will perform gateway monitoring by default which will generate two ICMP pings per second on the interface. Idle Timeout will not function in this case. This can be worked around by editing the gateway for this PPP link, and checking Disable Gateway Monitoring.

Compression (vjcomp) This option controls whether or not Van Jacobson TCP header compression will be used. By default it will be negotiated with the peer during login, so if both sides support the feature it will be used. Checking Disable vjcomp will cause the feature to always be disabled. Normally this feature is beneficial because it saves several bytes per TCP data packet. The option should almost always remain enabled. This compression is ineffective for TCP connections with enabled modern extensions like time stamping or SACK, which modify TCP options between sequential packets.

TCP MSS Fix The tcpmssfix option causes the PPP daemon to adjust incoming and outgoing TCP SYN segments so that the requested maximum segment size (MSS) is not greater than the amount allowed by the interface MTU. This is necessary in most cases to avoid problems caused by routers that drop ICMP “Datagram Too Big” messages. Without these messages, the originating machine sends data, it passes the rogue router then hits a machine that has an MTU that is not big enough for the data. Because the IP "Don't Fragment" option is set, this machine sends an ICMP "Datagram Too Big" message back to the originator and drops the packet. The rogue router drops the ICMP message and the originator never gets to discover that it must reduce the fragment size or drop the IP Don't Fragment option from its outgoing data. If this behavior is undesirable, check Disable tcpmssfix.

Note: The MTU and MSS values for the interface may also be adjusted on the interface's configuration page under the Interfaces menu, such as Interfaces > WAN.

Short Sequence (ShortSeq) This option is only meaningful if MLPPP is negotiated. It proscribes shorter multi-link fragment headers, saving two bytes on every frame. It is not necessary to disable this for connections that are not multi-link. If MLPPP is active and this feature must be disabled, check Disable shortseq.

Address Control Field Compression (AFCComp) This option only applies to asynchronous link types. It saves two bytes per frame. To disable this, check Disable ACF Compression.

Protocol Field Compression (ProtoComp) This option saves one byte per frame for most frames. To disable this, check Disable Protocol Compression.

5.6 GRE (Generic Routing Encapsulation)

Generic Routing Encapsulation (GRE) is a method of tunneling traffic between two endpoints without encryption. It can be used to route packets between two locations that are not directly connected, which do not require encryption. It can also be combined with a method of encryption that does not perform its own tunneling.

Note: The GRE protocol was originally designed by Cisco, and it is the default tunneling mode on many of their devices.

GRE tunnels can carry either IPv4, IPv6, or both types of traffic at the same time.

5.6.1 GRE Interface Settings

Parent interface The interface upon which the GRE tunnel will terminate. Often this will be WAN or a WAN-type connection.

Remote Address The address of the remote peer. This is the address where the GRE packets will be sent by this firewall; The routable external address at the other end of the tunnel.

Local IPv4/IPv6 Tunnel Address The internal IPv4 and IPv6 address for the end of the tunnel on this firewall. The firewall will use this address for its own traffic in the tunnel, and tunneled remote traffic would be sent to this address by the remote peer.

Remote IPv4/IPv6 Tunnel Address The IPv4 and IPv6 address used by the firewall inside the tunnel to reach the far side. Traffic destined for the other end of the tunnel must use this address as a gateway for routing purposes.

IPv4/IPv6 Tunnel Subnet The subnet mask for the GRE interface address.

Add Static Route When set, the firewall adds an explicit static route for the remote inner tunnel address/subnet via the local tunnel address. This can help with reaching the remote subnet in cases where other route table entries may select the wrong path to that destination.

Description A short description of this GRE tunnel for documentation purposes.

5.6.2 GRE Interface Management

To create or manage a GRE interface:

- Navigate to Interfaces > Assignments, GRE tab

Note: The items in this list are managed in the usual way. See *Managing Lists in the GUI*.

- Click ✚ Add to create a new GRE instance
- Complete the settings as described in *GRE Interface Settings*
- Click Save
- Navigate to Interfaces > Assignments
- Select the new GRE interface in the Available network ports list
- Click ✚ Add
- Note the name given to the new interface (e.g. OPT1)
- Navigate to Interfaces > <name> where <name> corresponds to the name of the GRE interface (e.g. OPT1)
- Check Enable interface
- Enter a new name for the interface in Description (optional)
- Click Save

Then use the interface as any other WAN-type interface. The firewall automatically creates a dynamic gateway for routing purposes. Depending on the use case, the interface may need NAT or firewall rules, static routes, and so on.

5.7 GIF (Generic tunnel InterFace)

A Generic Tunneling Interface (GIF) is similar to *GRE*; Both protocols are a means to tunnel traffic between two hosts without encryption. In addition to tunneling IPv4 or IPv6 directly, GIF may be used to tunnel IPv6 over IPv4 networks and vice versa. GIF tunnels are commonly used to obtain IPv6 connectivity to a tunnel broker such as Hurricane Electric in locations where IPv6 connectivity is unavailable.

See also:

See *Configuring IPv6 Through A Tunnel Broker Service* for information about connecting to a tunnel broker service.

GIF interfaces carry more information across the tunnel than can be done with GRE, but GIF is not as widely supported. For example, a GIF tunnel is capable of bridging layer 2 between two locations while GRE cannot.

GIF interfaces can carry IPv4 or IPv6 traffic, but not both at the same time.

5.7.1 GIF Interface Management

Parent interface The interface upon which the GIF tunnel will terminate. Often this will be WAN or a WAN-type connection.

GIF Remote Address The address of the remote peer. This is the address where the GIF packets will be sent by this firewall; The routable external address at the other end of the tunnel. For example, in a IPv6-in-IPv4 tunnel to Hurricane Electric, this would be the IPv4 address of the tunnel server, such as 209.51.181.2.

GIF tunnel local address The internal address for the end of the tunnel on this firewall. The firewall will use this address for its own traffic in the tunnel, and tunneled remote traffic would be sent to this address by the remote peer. For example, when tunneling IPv6-in-IPv4 via Hurricane Electric, they refer to this as the Client IPv6 Address.

GIF tunnel remote address The address used by the firewall inside the tunnel to reach the far side. Traffic destined for the other end of the tunnel must use this address as a gateway for routing purposes. For example, when tunneling IPv6-in-IPv4 via Hurricane Electric, they refer to this as the Server IPv6 Address.

GIF Tunnel Subnet The subnet mask or prefix length for the interface address. Typically 64.

ECN Friendly Behavior The ECN friendly behavior option controls whether or not the Explicit Congestion Notification (ECN)-friendly practice of copying the TOS bit into/out of the tunnel traffic is performed by the firewall. By default the firewall clears the TOS bit on the packets or sets it to 0, depending on the direction of the traffic. With this option set, the bit is copied as needed between the inner and outer packets to be more friendly with intermediate routers that can perform traffic shaping. This behavior breaks RFC 2893 so it must only be used when both peers agree to enable the option.

Outer Source Filtering When set, the firewall will not automatic filter based on the outer GIF source. This is normally desirable as it ensures a match with the configured remote peer, which is more secure. When disabled, martian and inbound filtering is not performed which allows asymmetric routing of the outer traffic. This is less secure, but some GIF peers may source traffic in this manner.

Description A short description of this GIF tunnel for documentation purposes.

5.7.2 GIF Interface Configuration

To create or manage a GIF interface:

- Navigate to Interfaces > Assignments, GIF tab

Note: The items in this list are managed in the usual way. See *Managing Lists in the GUI*.

- Click ✚ Add to create a new GIF instance
- Complete the settings as described in *GIF Interface Management*
- Click Save
- Navigate to Interfaces > Assignments
- Select the new GIF interface in the Available network ports list
- Click ✚ Add

10.7. GIF (Generic tunnel InterFace)

- Note the name given to the new interface (e.g. OPT1)
- Navigate to Interfaces > <name> where <name> corresponds to the name of the GIF interface (e.g. OPT1)
- Check Enable interface
- Enter a new name for the interface in Description (optional)
- Click Save

Then use the interface as any other WAN-type interface. The firewall automatically creates a dynamic gateway for routing purposes. Depending on the use case, the interface may need NAT or firewall rules, static routes, and so on.

5.8 LAGG (Link Aggregation)

Link aggregation is handled by lagg(4) type interfaces (LAGG) on AZTCO-FW® software. LAGG combines multiple physical interfaces together as one logical interface. There are several ways this can work, either for gaining extra bandwidth, redundancy, or some combination of the two.

Note: LACP will only work across multiple switches if the switches are Stackable.

To create or manage LAGG interfaces:

- Navigate to Interfaces > Assignments, LAGGs tab
- Click ➕ Add to create a new LAGG, or click ⧉ to edit an existing instance.
- Complete the settings as follows:

Parent Interfaces This list contains all currently unassigned interfaces, and members of the current LAGG interface when editing an existing instance. To add interfaces to this LAGG, select one or more interfaces in this list.

> Note: An interface may only be added to a LAGG group if it is not assigned. If an interface is not present in the list, it is likely already assigned as an interface.

LAGG Protocol There are currently six different operating modes for LAGG interfaces: LACP, Failover, Load Balance, Round Robin, and None.

> LACP The most commonly used LAGG protocol. This mode supports IEEE 802.3ad Link Aggregation Control Protocol (LACP) and the Marker Protocol. In LACP mode, negotiation is performed with the switch – which must also support LACP – to form a group of ports that are all active at the same time. This is knowns as a Link Aggregation Group, or LAG. The speed and MTU of each port in a LAG must be identical and the ports must also run at full- duplex. If link is lost to a port on the LAG, the LAG continues to function but at reduced capacity. In this way, an LACP LAGG bundle can gain both redundancy and increased bandwidth.
>
> Traffic is balanced between all ports on the LAG, however, for communication between two single hosts it will only use one single port at a time because the client will only talk to one MAC address at a time. For multiple connections through multiple devices, this limitation effectively becomes irrelevant. The limitation is also not relevant for failover.

10.8. LAGG (Link Aggregation)

In addition to configuring this option on AZTCO-FW, the switch must enable LACP on these ports or have the ports bundled into a LAG group. Both sides must agree on the configuration in order for it to work properly.

Failover When using the Failover LAGG protocol traffic will only be sent on the *primary* interface of the group. If the primary interface fails, then traffic will use the next available interface. The primary interface is the first interface selected in the list, and will continue in order until it reaches the end of the selected interfaces.

Note: By default, traffic may only be received on the active interface. Create a system tunable for net.link.lagg.failover_rx_all with a value of 1 to allow traffic to be received on every member interface.

Load Balance Load Balance mode accepts inbound traffic on any port of the LAGG group and balances outgoing traffic on any active ports in the LAGG group. It is a static setup that does not monitor the link state nor does it negotiate with the switch. Outbound traffic is load balanced based on all active ports in the LAGG using a hash computed using several factors, such as the source and destination IP address, MAC address, and VLAN tag.

Round Robin This mode accepts inbound traffic on any port of the LAGG group and sends outbound traffic using a round robin scheduling algorithm. Typically this means that traffic will be sent out in sequence, using each interface in the group in turn.

None This mode disables traffic on the LAGG interface without disabling the interface itself. The OS will still believe the interface is up and usable, but no traffic will be sent or received on the group.

Description A short note about the purpose of this LAGG instance.

- Click Save

After creating a LAGG interface, it works like any other physical interface. Assign the lagg interface under Interfaces > Assignments and give it an IP address, or build other things on top of it such as VLANs.

5.8.1 LAGG and Traffic Shaping

Due to limitations in FreeBSD, lagg(4) does not support altq(4) so it is not possible to use the traffic shaper on LAGG interfaces directly. vlan(4) interfaces support altq(4) and VLANs can be used on top of LAGG interfaces, so using VLANs can work around the problem. As an alternate workaround, Limiters can control bandwidth usage on LAGG interfaces.

5.8.2 LAGG Throughput

Using a LAGG does not necessarily guarantee full throughput equal to the sum of all interfaces. In particular, a single flow will not exceed the throughput of a LAGG member interface. Traffic on a LAGG is hashed in such a way that flows between two hosts, such as AZTCO-FW and an upstream gateway, would only use a single link since the flow is between a single MAC address on each side.

In networks where there are many hosts communicating with different MAC addresses, the usage can approach the sum of all interfaces in the LAGG.

10.8. LAGG (Link Aggregation)

5.9 AZTCO-FW QinQ Configuration

QinQ, also known as IEEE 802.1ad or stacked VLANs, is a means of nesting VLAN tagged traffic inside of packets that are already VLAN tagged, or "double tagging" the traffic.

QinQ is used to move groups of VLANs over a single link containing one outer tag, as can be found on some ISP, Metro Ethernet, or datacenter links between locations. It can be a quick/easy way of trunking VLANs across locations without having a trunking-capable connection between the sites, provided the infrastructure between the locations does not strip tags from the packets.

Setting up QinQ interfaces in AZTCO-FW® software is fairly simple:

- Navigate to Interfaces > Assignments
- Click the QinQ tab
- Click + Add to add a new QinQ entry
- Configure the QinQ entry as follows:

  Parent Interface The interface that will carry the QinQ traffic.

  First level tag The outer VLAN ID on the QinQ interface, or the VLAN ID given by the provider for the site-to-site link.

  Adds interface to QinQ interface groups When checked, a new interface group will be created called QinQ that can be used to filter all of the QinQ subinterfaces at once.

  When hundreds or potentially thousands of QinQ tags are present, this greatly reduces the amount of work needed to use the QinQ interfaces

  Description Optional text for reference, used to identify the entry

  Member(s) Member VLAN IDs for QinQ tagging. These can be entered one per row by clicking + Add Tag, or in ranges such as 100-150

- Click Save to complete the interface

In the following example (Figure *QinQ Basic Example*), a QinQ interface is configured to carry tagged traffic for VLANs *10* and *20* across the link on *igb3* with a first level tag of *2000*.

In Figure *QinQ List*, this entry is shown on the QinQ tab summary list.

The automatic interface group, shown in Figure *QinQ Interface Group*, must not be manually edited. Because these interfaces are not assigned, it is not possible to make alterations to the group without breaking it. To re-create the group, delete it from this list and then edit and save the QinQ instance again to add it back.

Rules may be added to the QinQ tab under Firewall > Rules to pass traffic in both directions across the QinQ links.

From here, how the QinQ interfaces are used is mostly up to the needs of the network. Most likely, the resulting interfaces may be assigned and then configured in some way, or bridged to their local equivalent VLANs (e.g. bridge an assigned igb2_vlan10 to igb3_2000_10 and so on).

The QinQ configuration will be roughly the same on both ends of the setup. For example, if both sides use identical interface configurations, then traffic that leaves Site A out on igb3_2000_10 will go through VLAN 2000 on igb3, come out the other side on VLAN 2000 on igb3 at Site B, and then in igb3_2000_10 at Site B.

See also:

- *Virtual LANs (VLANs)*

10.9. AZTCO-FW QinQ Configuration

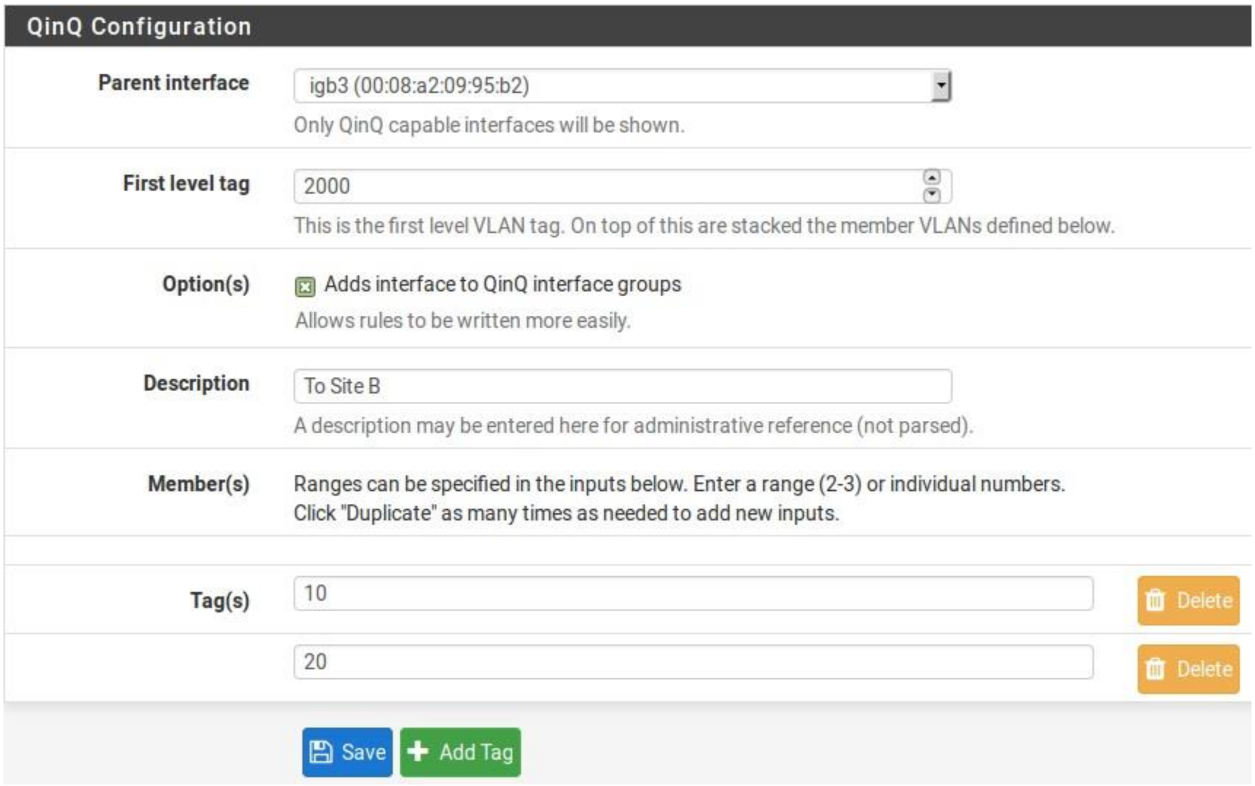

QinQ Configuration

Parent interface: igb3 (00:08:a2:09:95:b2)
Only QinQ capable interfaces will be shown.

First level tag: 2000
This is the first level VLAN tag. On top of this are stacked the member VLANs defined below.

Option(s): ☒ Adds interface to QinQ interface groups
Allows rules to be written more easily.

Description: To Site B
A description may be entered here for administrative reference (not parsed).

Member(s): Ranges can be specified in the inputs below. Enter a range (2-3) or individual numbers.
Click "Duplicate" as many times as needed to add new inputs.

Tag(s): 10
20

Fig. 3: QinQ Basic Example

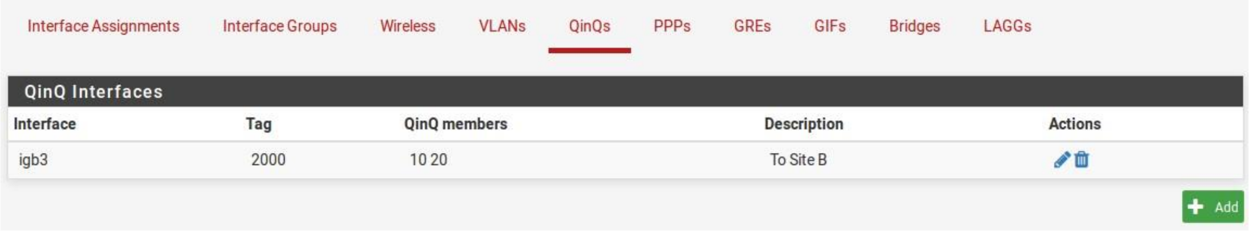

QinQ Interfaces

Interface	Tag	QinQ members	Description	Actions
igb3	2000	10 20	To Site B	

Fig. 4: QinQ List

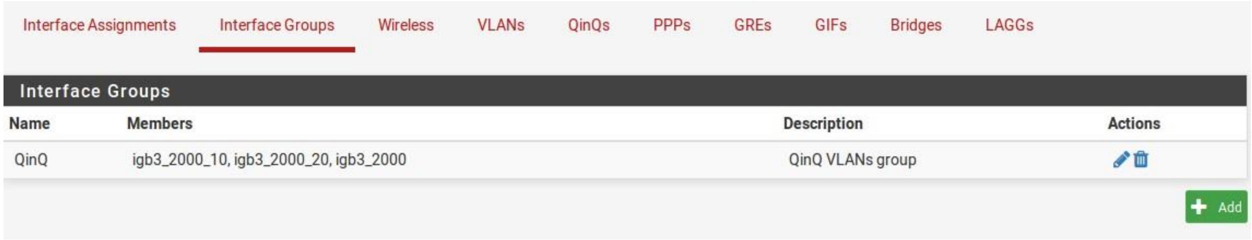

Interface Groups

Name	Members	Description	Actions
QinQ	igb3_2000_10, igb3_2000_20, igb3_2000	QinQ VLANs group	

Fig. 5: QinQ Interface Group

10.9. AZTCO-FW QinQ Configuration

AZTCO-FW® software supports numerous types of network interfaces, either using physical interfaces directly or by employing other protocols such as PPP or VLANs.

Interface assignments and the creation of new virtual interfaces are all handled under Interfaces > Assignments.

5.10 Physical and Virtual Interfaces

Most interfaces discussed in this chapter can be assigned as WAN, LAN, or an OPT interface under Interfaces > Assignments. All currently-defined and detected interfaces are listed directly on Interfaces > Assignments or in the list of interfaces available for assignment. By default, this list includes only the physical interfaces, but the other tabs under Interfaces > Assignments can create virtual interfaces which can then be assigned.

Interfaces on AZTCO-FW support various combinations of options on the interfaces themselves. They can also support multiple networks and protocols on a single interface, or multiple interfaces can be bound together into a larger capacity or redundant virtual interface.

All interfaces are treated equally; Every interface can be configured for any type of connectivity or role. The default WAN and LAN interfaces can be renamed and used in other ways.

Physical interfaces and virtual interfaces are treated the same once assigned, and have the same capabilities. For example, a VLAN interface can have the same type of configuration that a physical interface can have. Some interface types receive special handling once assigned, which are covered in their respective sections of this chapter.

This section covers the various types of interfaces that can be created, assigned, and managed.

5.11 Limitations

While AZTCO-FW software does not impose any limits on the number of interfaces, large numbers of interfaces may function in suboptimal ways. For example, the firewall may take much longer configure the interfaces, and the GUI may have rendering issues with large numbers of tabs or menu entries.

Most hardware will accommodate as many physical interfaces as can fit into the case. Issues may vary from driver to driver but generally are hardware-related and not the result of the operating system or AZTCO-FW software.

Note: With a large number of physical interfaces, the number of mbufs will likely need to be increased. See *Tuning and Troubleshooting Network Cards*.

Physical limitations aside, AZTCO-FW software can also handle significant numbers of virtual interfaces such as VLANs, LAGGs, VPNs, and more. These types interfaces tend to outnumber physical interfaces, especially VLANs.

Issues reported by users with large numbers of interfaces (physical and virtual) vary by hardware, configuration, and browser. These issues tend to increase as the number of interfaces approaches 200. Should a particular environment require more than 128 interfaces, consider alternate designs that do not involve using all of the interfaces on the firewall directly. If the firewall must handle large numbers of interfaces, be wary of potential performance and GUI concerns.

CHAPTER

SIX

USER MANAGEMENT AND AUTHENTICATION

6.1 AZTCO-FW Default Username and Password

The default credentials for a AZTCO-FW® firewall are:

- Username: *admin*
- Password: *aztco*

6.2 Privileges

Managing privileges for users and groups is done similarly, so both will be covered here rather than duplicating the effort. Whether a user or group is managed, the entry must be created and saved first before privileges can be added to the account or group. To add privileges, when editing the existing user or group, click + Add in the Assigned Privileges or Effective Privileges section.

A list of all available privileges is presented. Privileges may be added one at a time by selecting a single entry, or by multi-select using ctrl-click. If other privileges are already present on the user or group, they are hidden from this list so they cannot be added twice. To search for a specific privilege by name, enter the search term in the Filter box and

click ▼ Filter.

Selecting a privilege will show a short description of its purpose in the information block area under the permission list and action buttons. Most of the privileges are self-explanatory based on their names, but a few notable permissions are:

WebCfg - All Pages Lets the user access any page in the GUI

WebCfg - Dashboard (all) Lets the user access the dashboard page and all of its associated functions (widgets, graphs, etc.)

WebCfg - System User Password Manager Page: If the user has access to only this page, they can login to the GUI to set their own password but do nothing else.

User - VPN - IPsec xauth Dialin Allows the user to connect and authenticate for IPsec xauth

User - Config - Deny Config Write Does not allow the user to make changes to the firewall config (*config.xml*). Note that this does not prevent the user from taking other actions that do not involve writing to the config.

User - System - Shell account access Gives the user the ability to login over ssh, though the user will not have root-level access so functionality is limited. A package for *sudo* is available to enhance this feature.

After login, the firewall will attempt to display the dashboard. If the user does not have access to the dashboard, they will be forwarded to the first page in their privilege list which they have permission to access.

Menus on the firewall only contain entries for which privileges exist on a user account. For example, if the only Diagnostics page that a user has access to is Diagnostics > Ping then no other items will be displayed in the Diagnostics menu.

6.3 Manage Local Users

The Users tab under System > User Manager is where individual users are managed. To add a new user, click + Add, to edit an existing user, click ⧉.

Note: The admin user cannot be deleted and its username may not be changed.

Before permissions may be added to a user, it must first be created, so the first step is always to add the user and save. If multiple users need the same permissions, it is easier to add a group and then add users to the group.

To add a user, click + Add and the new user screen will appear.

Disabled This checkbox controls whether this user will be active. If this account should be deactivated, check this box.

Username Sets the login name for the user. This field is required, must be 16 characters or less and may only contain letters, numbers, and a period, hyphen, or underscore.

Password and Confirmation are also required. Passwords are stored in the AZTCO-FW® configuration as hashes. Ensure the two fields match to confirm the password.

Full Name Optional field which can be used to enter a longer name or a description for a user account.

Expiration Date May also be defined if desired to deactivate the user automatically when that date has been reached. The date must be entered in *MM/DD/YYYY* format.

Group Memberships If groups have already been defined (*Manage Local Groups*), this control may be used to add the user as a member. To add a group for this user, find it in the Not Member Of column, select it, and click » to move it to the Member Of column. To remove a user from the group, select it from the Member Of column and click « to move it to the Not Member Of column.

Effective Privileges Appears when editing an existing user, not when adding a user. See *Privileges* for information on managing privileges. If the user is part of a group, the group's permissions are shown in this list but those permissions cannot be edited, however additional permissions may be added.

Certificate Behavior of this section changes depending on whether a user is being added or edited. When adding a user, to create a certificate check Click to create a user certificate to show the form to create a certificate. Fill in the Descriptive name, choose a Certificate Authority, select a Key Length, and enter a Lifetime. For more information on these parameters, see *Create an Internal Certificate*. If editing a user, this section of the page instead becomes a list of user certificates. From

here, click + Add to add a certificate to the user. The settings on that page are identical to

11.3. Manage Local Users

Create an Internal Certificate except even more of the data is pre- filled with the username. If the certificate already exists, select *Choose an Existing Certificate* and then pick an Existing Certificate from the list.

Authorized keys SSH public keys may be entered for shell or other SSH access. To add a key, paste or enter in the key data.

IPsec Pre-Shared Key Used for a non-xauth Pre-Shared Key mobile IPsec setup. If an IPsec PreShared Key is entered here, the username is used as the identifier. The PSK is also displayed under VPN > IPsec on the Pre- Shared Keys tab. If mobile IPsec will only be used with xauth, this field may be left blank.

After saving the user, click ⧉ on the user's row to edit the entry if necessary.

6.1.1 Per-user GUI Options and Dashboard Layout

Each user can have their own settings for various GUI options and their dashboard layout. To enable this for a user, check the Custom Settings box when adding or editing the user. The user then automatically gets their own dashboard layout, starting from the system-wide layout. Choose the other GUI options desired for the user such as theme, top navigation, host name in menu, dashboard columns, show/hide associated panels, left column labels and browser tab text.

Tip: Users who want to adjust their own GUI options need the WebCfg - System: User Settings privilege.

Users in the admin group already have this privilege.

A user with Custom Settings enabled (and the User Settings privilege) will have menu option System > User Settings. The user can select this to change the desired GUI options for their user name.

When a user with Custom Settings enabled adds, moves or removes dashboard widgets, the custom dashboard layout is saved in the preferences for only that user.

6.4 Manage Local Groups

Groups are a great way to manage sets of permissions to give users so that they do not need to be maintained individually on every user account. For example, a group could be used for IPsec xauth users, or a group that can access the firewall's dashboard, a group of firewall administrators, or many other possible scenarios using any combination of privileges.

As with users, a group must first be created before privileges can be added. After saving the group, edit the group to add privileges.

Groups are managed under System > User Manager on the Groups tab. To add a new group from this screen, click + Add. To edit an existing group, click ⧉ next to its entry in the list.

Note: When working with LDAP and RADIUS, local groups must exist to match the groups the users are members of on the server. For example, if an LDAP group named "firewall_admins" exists then AZTCO-FW must also contain a group named identically, "firewall_admins", with the desired privileges. Remote groups with long names or names containing spaces or other special characters must be configured for a *Remote* Scope.

11.4. Manage Local Groups

Start the process of adding a group by clicking + Add and the screen to add a new group will appear.

Group name This setting has the same restrictions as a username: It must be 16 characters or less and may only contain letters, numbers, and a period, hyphen, or underscore. This can feel somewhat limited when working with groups from LDAP, for example, but usually it's easier to create or rename an appropriately-named group on the authentication server instead of attempting to make the firewall group match.

Scope Can be set *Local* for groups on the firewall itself (such as those for use in the shell), or *Remote* to relax the group name restrictions and to prevent the group name from being exposed to the base operating system. For example, *Remote* scope group names may be longer, and may contain spaces.

Description Optional free-form text for reference and to better identify the purpose of the group in case the Group name is not sufficient.

Group Memberships This set of controls defies which existing users will be members of the new group. Firewall users are listed in the Not Members column by default. To add a user to this group, find

it in the Not Members column, select it, and click » to move it to the Members column. To

remove a user from the group, select it from the Members column and click « to move it to the Not Members column.

Assigned Privileges Appears only when editing an existing group. This section allows adding privileges to the group. See *Privileges* earlier in this for information on managing privileges.

6.5 Authentication Servers

AZTCO-FW® software can use RADIUS and LDAP servers to authenticate users from remote sources. *Support Throughout AZTCO-FW* contains information on which areas of the firewall support these servers

To add a new server:

- Navigate to System > User Manager, Authentication Servers tab
- Click ➕ Add

To edit an existing server, click ⧉ next to its entry on the same page.

See also:

AZTCO-FW Hangouts on Youtube to view the August 2015 Hangout on RADIUS and LDAP.

Each type of authentication server is covered in the following documents

5.1.1 RADIUS Authentication Servers

Remote Authentication Dial-In User Service (RADIUS) is a protocol commonly supported by a wide variety of networking equipment for user authentication, authorization, and accounting (AAA).

Servers are commonly available as well, including *FreeRADIUS* and *Active Directory via NPS*.

Though most areas on AZTCO-FW® software which support RADIUS now integrate their RADIUS settings via the user manager, a few remain which use separate settings, such as the PPPoE and L2TP servers.

See also:

- *Controlling Client Parameters via RADIUS*

> Warning: Secure the link between the firewall and the RADIUS server. If the server is local, use a trusted management network. If the server is remote, communicate only over VPN tunnels.
>
> Some RADIUS protocols transmit passwords in plain text, and though others attempt to protect the password in other ways, other aspects of the protocol are not encrypted and may contain sensitive information.

RADIUS Configuration

Descriptive name The name for this RADIUS server. This name will be used to identify the server throughout the GUI.

Protocol The protocol used by the firewall when performing RADIUS requests. May be one of:

PAP Password Authentication Protocol. Sends passwords unencrypted, and is considered weak. It is more widely supported than other methods, and may be required by specific features (e.g. mOTP).

> Warning: Due to its security deficiencies, avoid using PAP where possible.

MD5-CHAP Challenge-Handshake Authentication Protocol using MD5 hashing. The RADIUS server sends a challenge value and the client responds with a hash of the challenge value and the password together. More secure than PAP as it does not transmit passwords in the clear, but both parties must know the plain text of the password.

MS-CHAPv1 A Microsoft variation of CHAP where neither side needs to know the plain text of the password. Though it is generally more secure, it has other known weaknesses which make it vulnerable to attack.

MS-CHAPv2 An updated variation of MS-CHAPv1. It is used in EAP as well as 802.1x/WPA Enterprise for wireless. However, it also has known weaknesses.

Note: Certain RADIUS features may require specific modes. For example, mOTP typically requires PAP since it reads the password in the clear to separate the PIN and OTP code. Services utilizing EAP typically use MS-CHAPv2.

Hostname or IP address The address of the RADIUS server. This can be a fully qualified domain name or an IPv4 IP address.

Warning: The RADIUS client on the firewall does not currently support IPv6.

Shared Secret The password established for this firewall *on the RADIUS server* software.

Services offered This selector sets which services are offered by this RADIUS server.

Authentication The firewall will use this RADIUS server to authenticate users.

Accounting The firewall will send RADIUS start/stop accounting packet data for login sessions if supported in the area where it is used.

Authentication and Accounting The server will be used for both types of actions.

Authentication port Only appears if an Authentication mode is chosen. Sets the UDP port where RADIUS authentication will occur. The default RADIUS authentication port is 1812.

Accounting port Only appears if an Accounting mode is chosen. Sets the UDP port where RADIUS accounting will occur. The default RADIUS accounting port is 1813.

Authentication Timeout Controls how long, in seconds, that the RADIUS server may take to respond to an authentication request. If left blank, the default value is 5 seconds. If an interactive two-factor authentication system is in use, increase this timeout to account for how long it will take the user to receive and enter a token, which can be 60-120 seconds or more if it must wait for an external action such as a phone call, SMS message, etc.

RADIUS NAS IP Attribute Sets the value the firewall will send in the RADIUS request NAS-IP-Address attribute. This value is used by the RADIUS server to identify this firewall. The server can use this value to make authentication decisions, or to denote which node users were authenticated by in accounting data.

In most cases, the NAS-IP-Address value does not matter so long as it is unique to this firewall. However, more complicated RADIUS environments may use this attribute to let the server make more informed decisions about users logging into different services. For example, if there are multiple Captive Portal instances on the firewall, multiple RADIUS server entries can be created, each using the specific interface address for a given portal. The RADIUS server could then choose to only let certain sets of users login to each portal.

Adding a RADIUS Server

To add a new RADIUS server:

- Add the firewall as a client on the RADIUS server
- Navigate to System > User Manager, Authentication Servers tab
- Click + Add
- Set the Type selector to *RADIUS*

  The GUI will change the form to display RADIUS Server Settings
- Fill in the fields as described in *RADIUS Configuration*
- Click Save to create the server
- Navigate to Diagnostics > Authentication to test the RADIUS server using a valid account.

RADIUS Groups

There are two requirements for RADIUS groups to function properly:

- The RADIUS server must return a list of groups in the Class RADIUS reply attribute as a string.
- The same groups must exist locally (*Manage Local Groups*)

Multiple groups groups returned by the RADIUS server in the Class attribute must be separated by a semicolon. For example, in FreeRADIUS, to return the admins and VPNUsers groups, use the following Reply-Item RADIUS Attribute:

```
Class := "admins;VPNUsers"
```

If the RADIUS server returns the group list properly for a user, and the groups exist locally, then the groups will be listed on the results when using the Diagnostics > Authentication page to test an account.

If the groups do not show up when testing, ensure the groups exist in the *Group Manager* with matching names and that the server is returning the Class attribute as a string, not binary.

6.1.2 LDAP Authentication Servers

Though Lightweight Directory Access Protocol (LDAP) is technically a repository for user information, it also supports mechanisms for user authentication via bind operations.

There are many popular user directory implementations which use LDAP, including Active Directory, OpenLDAP, FreeIPA, and more.

Note: LDAP server implementations and schemas vary widely. As such, there are no complete and specific examples in this document.

LDAP Configuration

Hostname or IP address The address of the LDAP server. This can be a fully qualified domain name, an IPv4 IP address, or an IPv6 IP address.

Note: If this LDAP server uses SSL, the value of this field must match the certificate presented by the LDAP server. Typically this means it must be a hostname which resolves to the IP address of the LDAP server, but the specific requirements depend on the contents of the server certificate.

For example, with a value of ldap.example.com in this field, the server certificate must include an FQDN value of ldap.example.com, and ldap.example.com must resolve to 192. 168.1.5. One exception to this is if the IP address of the server also happens to be the listed in the server certificate.

This can be worked around in some cases by creating a DNS host override to make the server certificate hostname resolve to the correct IP address if they do not match in this network infrastructure and they cannot be easily fixed.

Port value This setting specifies the port on which the LDAP server is listening for LDAP queries. The default port is 389 for Standard TCP and STARTTLS, and 636 for SSL. This field is updated automatically with the proper default value based on the selected Transport.

Note: When using port 636 for SSL, AZTCO-FW® uses an ldaps:// URL, not STARTTLS. Ensure that the LDAP server is listening on the correct port with the correct mode.

Transport This setting controls which transport method will be used by the firewall to communicate with the LDAP server.

Warning: LDAP queries will contain sensitive data, such as usernames, passwords, and other information about the user. The best practice is for the firewall to use encryption when communicating with the LDAP server, if the LDAP server supports it. Both SSL/TLS and STARTTLS will encrypt traffic between the firewall and the LDAP server.

Standard TCP (Default) Plain unencrypted TCP connections on port 389. This is not secure, but is widely supported and also useful for debugging with packet captures. Do not use this protocol across untrusted networks.

STARTTLS Encrypted Connects using TCP port 389 but negotiates encryption with the server using STARTTLS.

Note: Not all LDAP servers support STARTTLS, check the LDAP server documentation and configuration.

SSL/TLS Encrypted Connects using SSL/TLS on TCP port 636 to encrypt LDAP queries.

Note: Not all LDAP servers support SSL/TLS, check the LDAP server documentation and configuration.

Peer Certificate Authority The CA chosen with this selector is used by the firewall to validate the LDAP server certificate when Transport is set to SSL/TLS Encrypted or STARTTLS Encrypted mode.

The selected CA must match the CA which signed the LDAP server certificate, otherwise validation will fail. If the LDAP server is using a globally trusted certificate (e.g. Let's Encrypt or another public CA), choose *Global Root CA List*.

See *Certificate Authority Management* for more information on creating or importing CAs.

Client Certificate (Factory only) This certificate is sent to the LDAP server to identify this client when using an encrypted transport mode. If the LDAP server requires a client certificate, the server will use this certificate to ensure that the firewall is authorized to make LDAP queries.

This certificate must be issued by the CA used by the LDAP server to validate connecting clients.

Protocol version Chooses which version of the LDAP protocol is employed by the LDAP server, either *2* or *3*, typically *3*.

Server Timeout The time, in seconds, after which LDAP operations are considered as failed. Using a lower value will allow the GUI to try other authentication sources faster when the server fails. If the LDAP server is slow or overloaded, a larger value can help the firewall accept delayed responses.

Search scope Determines where, and how deep, an LDAP search will be performed to locate a match.

Level Controls the depth of the LDAP search.

One Level Search only one level, defined by the Authentication Containers.

Entire Subtree Search the entire subtree of the directory, starting with the Authentication Containers.

Tip: This is typically the best choice, and is nearly always required for Active Directory configurations.

Base DN Controls where the search will start. Typically set to the root of the LDAP structure, e.g. DC=example,DC=com

Authentication containers A list of potential account locations or containers, separated by semicolons. These containers will be prepended to the Base DN above when the firewall crafts LDAP queries. Alternately, specify a full container path here and leave the Base DN blank.

Tip: If the LDAP server supports it, and the bind settings are correct, click 🔍 Select a container to browse the LDAP server and select containers from a list.

Some examples of containers are:

- CN=Users;DC=example;DC=com This searches for users inside of the domain component example.com, a common syntax for Active Directory
- CN=Users,DC=example,DC=com;OU=OtherUsers,DC=example,DC=com This searches in two different locations, the second of which is restricted to the OtherUsers organizational unit.

Extended Query Specifies an extra restriction to query after the username, which allows group membership to be used as a filter. This must include both the item to search as well as the method of searching. For example, a restriction based on group membership would use memberOf. Check the LDAP server documentation for information on forming such queries.

To set an extended query, check the box and fill in the Query value with a filter such as:

```
memberOf=CN=VPNUsers,CN=Users,DC=example,DC=com
```

Bind credentials Controls how this LDAP client will attempt to bind to the server.

Note: Active Directory typically requires the use of bind credentials and may need a service account or administrator-equivalent depending on the server configuration. Consult Windows documentation to determine which is necessary in a specific environment.

Bind Anonymous (Default) When checked the firewall will use anonymous binds. When unchecked the GUI presents the Bind Credentials fields.

Bind Credentials (User DN/Password) When Bind Anonymous is unchecked, the credentials in these fields are used by the firewall to make authenticated binds when performing a query.

The User DN may be a username or a full DN, depending on what the LDAP server requires.

Attributes

Initial Template This option only appears when initially creating an LDAP server entry. It pre-fills the remaining options on the page with common defaults for a given type of LDAP server. The choices include *OpenLDAP*, *Microsoft AD*, and *Novell eDirectory*.

User naming attribute The attribute used to identify the name of a user, most commonly cn or samAccountName.

Group naming attribute The attribute used to identify a group, such as cn.

Group member attribute The attribute of a user that signifies it is the member of a group, such as member, memberUid, memberOf, or uniqueMember.

RFC2307 Groups Specifies how group membership is organized on the LDAP server. When unset (default), the queries assume the server uses Active Directory style group membership (RFC 2307bis) where groups are listed as an attribute of the user object. When checked, queries use RFC 2307 style group membership where the users are listed as members on the group object.

Note: In this mode the Group member attribute will typically be set to memberUid, but may vary by LDAP schema.

RFC2307 User DN When set, queries include the user DN when searching for groups.

Group Object Class Specifies the object class of RFC 2307 style groups. Typically posixGroup but it may vary by LDAP schema. Not necessary for Active Directory style groups.

UTF8 Encode When checked, queries to the LDAP server are encoded for UTF-8 and the responses are decoded from UTF-8. Support varies depending on the LDAP server. Generally only necessary if user names, groups, passwords, and other attributes contain UTF-8 or international style accented characters.

Username Alterations When unchecked, a username given as user@hostname will have the @hostname portion stripped so only the username is sent in the LDAP bind request. When checked, the username is sent in full.

Allow Unauthenticated Bind When set, bind requests with empty passwords will be rejected locally. Some LDAP servers, specifically Microsoft Active Directory, will accept unauthenticated bind requests and treat them as successful.

> Warning: This behavior must be disabled on the LDAP server where possible. Allowing requests to succeed with an empty password is a significant security risk and it affects any device or service authenticating against an LDAP server.
>
> Though this option allows AZTCO-FW software to reject such authentication attempts, other LDAP clients may not offer the same choice. Disabling the feature on the server is the most secure means of correcting the problem. Consult the LDAP server documentation for information on disabling this behavior.

Adding an LDAP Server

To add a new LDAP server:

- Make sure that the LDAP server can be reached by the firewall
- Import the Certificate Authority used by the LDAP server into AZTCO-FW before proceeding if using SSL/TLS or
  STARTTLS encryption

  See *Certificate Authority Management* for more information on creating or importing CAs.
- Navigate to System > User Manager, Authentication Servers tab

- Click + Add
- Set the Type selector to *LDAP*

  The GUI will change the form to display LDAP server settings
- Fill in the fields as described previously in *LDAP Configuration*
- Click Save to create the server
- Visit Diagnostics > Authentication to test the LDAP server using a valid account

LDAP Groups

There are two requirements for LDAP groups to function properly:

- The LDAP authentication settings must match the group membership style used by the LDAP server
- The same groups must exist locally (*Manage Local Groups*)

If the LDAP query returns the group list properly for a user, and the groups exist locally, then the groups will be listed on the results when using the Diagnostics > Authentication page to test an account.

If the groups do not show up, ensure they exist on AZTCO-FW with matching names and that the proper group structure is present on the LDAP authentication server entry (e.g. RFC 2703 options.)

6.6 Settings

The Settings tab in the User Manager controls two things: How long a login session is valid, and where the GUI logins will prefer to be authenticated.

Session Timeout This field specifies how long a GUI login session will last when *idle*. This value is specified in minutes, and the default is four hours (240 minutes). A value of 0 may be entered to disable session expiration, making the login sessions valid forever. A shorter timeout is better, though make it long enough that an active administrator would not be logged out unintentionally while making changes.

> Warning: Allowing a session to stay valid when idle for long periods of time is insecure. If an administrator leaves a terminal unattended with a browser window open and logged in, someone or something else could take advantage of the open session.

Authentication Server This selector chooses the primary authentication source for users logging into the GUI. This can be a RADIUS or LDAP server, or the default *Local Database* . If the RADIUS or LDAP server is unreachable for some reason, the authentication will fall back to *Local Database* even if another method is chosen.

When using a RADIUS or LDAP server, the users and/or group memberships must still be defined in the firewall in order to properly allocate permissions, as there is not yet a method to obtain permissions dynamically from an authentication server.

For group membership to work properly, AZTCO-FW must be able to recognize the groups as presented by the authentication server. This requires two things:

1. The local groups must exist with identical names (*Manage Local Groups*).

2. AZTCO-FW must be able to locate or receive a list of groups from the authentication server.

11.6. Settings
See *Authentication Servers* for details specific to each type of authentication server.

6.7 Logging Out of the webGUI

In current versions of AZTCO-FW® software, log off by navigating to System > Logout or by closing the browser window.

Sessions will automatically expire if they are idle for longer than the Session Timeout defined on System > User Manager, Settings tab. The default session timeout is 4 hours (240 minutes) of idle time.

See also:

Sudo Package External User Authentication Examples Granting Users Access to SSH Accessing the Firewall Filesystem with SCP Authenticating Users with Google Cloud Identity Troubleshooting Authentication Troubleshooting Access when Locked Out of the Firewall

The User Manager in AZTCO-FW® software provides the ability to create and manage multiple user accounts. These accounts can be used to access the GUI, use VPN services like OpenVPN and IPsec, and use the Captive Portal.

The User Manager is located at System > User Manager. From there users, groups, servers may be managed, and settings that govern the behavior of the User Manager may be changed.

The User Manager can also be used to define external authentication sources such as RADIUS and LDAP.

See also:

AZTCO-FW Hangouts on Youtube to view the February 2015 Hangout on User Management and Privileges, and the August 2015 Hangout on RADIUS and LDAP.

6.8 Support Throughout AZTCO-FW

As of this writing, not all areas of AZTCO-FW hook back into the User Manager.

AZTCO-FW GUI Supports users in the User Manager, and via RADIUS or LDAP. Groups or Users from RADIUS or LDAP require definitions in the local User Manager to manage their access permissions. OpenVPN Supports users in the User Manager, RADIUS or LDAP via User Manager.

IPsec Supports users in the User Manager, RADIUS or LDAP via User Manager for Xauth, and RADIUS for IKEv2 with EAP-RADIUS.

Captive Portal Support local users in the User Manager, and RADIUS users via settings in the Captive Portal page.

L2TP Supports users in the L2TP settings, and via RADIUS in the L2TP settings.

PPPoE Server Supports users in the PPPoE settings, and via RADIUS in the PPPoE settings.

CHAPTER

SEVEN

CERTIFICATE MANAGEMENT

7.1 Certificate Authority Management

Certificate Authorities (CAs) are managed from System > Cert Manager, on the CAs tab. From this screen CAs may be added, edited, exported, or deleted.

7.1.1 Create a new Certificate Authority

To create a new CA, start the process as follows:

- Navigate to System > Cert Manager on the CAs tab.
- Click Add to create a new a CA.
- Enter a Descriptive name for the CA. This is used as a label for this CA throughout the GUI.
- Select the Method that best suits how the CA will be generated. These options and further instructions are in the corresponding sections below:
  - Create an Internal Certificate Authority
  - Import an Existing Certificate Authority
  - Create an Intermediate Certificate Authority

Create an Internal Certificate Authority

The most common Method used from here is to *Create an Internal Certificate Authority*. This will make a new root CA based on information entered on this page.

- Select the Key length to choose how "strong" the CA is in terms of encryption. The longer the key, the more secure it is. However, longer keys can take more CPU time to process, so it isn't always wise to use the maximum value. The default value of *2048* is a good balance.
- Select a Digest Algorithm from the supplied list. The current best practice is to use an algorithm stronger than SHA1 where possible. *SHA256* is a good choice.

  Note: Some older or less sophisticated equipment, such as VPN-enabled VoIP handsets may only support SHA1 for the Digest Algorithm. Consult device documentation for specifics.

- Enter a value for Lifetime to specify the number of days for which the CA will be valid. The duration depends on personal preferences and site policies. Changing the CA frequently is more secure, but it is also a management headache as it would require reissuing new certificates when the CA expires. By default the GUI suggests using 3650 days, which is approximately 10 years.
- Enter values for the Distinguished name section for personalized parameters in the CA. These are typically filled in with an organization's information, or in the case of an individual, personal information. This information is mostly cosmetic, and used to verify the accuracy of the CA, and to distinguish one CA from another. Punctuation and special characters must not be used.
  - Select the Country Code from the list. This is the ISO-recognized country code, not a hostname top-level domain.
  - Enter the State or Province fully spelled out, not abbreviated.
  - Enter the City.
  - Enter the Organization name, typically the company name.
  - Enter a valid Email Address.
  - Enter the Common Name (CN). This field is the internal name that identifies the CA. Unlike a certificate, the CN for a CA does not need to be the hostname, or anything specific. For instance, it could be called *VPNCA* or *MyCA*.

    Note: Although it is technically valid, avoid using spaces in the CN.
- Click Save

If errors are reported, such as invalid characters or other input problems, they will be described on the screen. Correct the errors, and attempt to Save again.

Import an Existing Certificate Authority

If an existing CA from an external source needs to be imported, it can be done by selecting the Method of *Import an Existing Certificate Authority*. This can be useful in two ways: One, for CAs made using another system, and two, for CAs made by others that must be trusted.

- Enter the Certificate data for the CA. To trust a CA from another source, only the Certificate data for the CA is required. It is typically contained in a file ending with *.crt* or *.pem*. It would be plain text, and enclosed in a block such as:

```
-----BEGIN CERTIFICATE-----
[A bunch of random-looking base64-encoded data]
-----END CERTIFICATE-----
```

- Enter the Certificate Private Key if importing a custom external CA, or a CA that is capable of generating its own certificates and certificate revocation lists. This is typically in a file ending in *.key*. It would be plain text data enclosed in a block such as:

```
-----BEGIN RSA PRIVATE KEY----[A bunch of random-looking base64-encoded data]
-----END RSA PRIVATE KEY-----
```

- Enter the Serial for next certificate if the private key was entered. This is essential. A CA will create certificates each with a unique serial number in sequence. This value controls what the serial will be for the next certificate generated from this CA. It is essential that each certificate have a unique serial, or there will be problems later with certificate revocation. If the next serial is unknown, attempt to estimate how many certificates have been made from the CA, and then set the number high enough a collision would be unlikely.

- Click Save

If errors are reported, such as invalid characters or other input problems, they will be described on the screen. Correct the errors, and attempt to Save again.

Importing a Chained or Nested Certificate Authority

If the CA has been signed by an intermediary and not directly by a root CA, it may be necessary to import both the root and the intermediate CA together in one entry, such as:

```
-----BEGIN CERTIFICATE-----
[Subordinate/Intermediate CA certificate text]
-----END CERTIFICATE---------BEGIN CERTIFICATE-----
[Root CA certificate text]
-----END CERTIFICATE-----
```

Create an Intermediate Certificate Authority

An Intermediate CA will create a new CA that is capable of generating certificates, yet depends on another CA higher above it. To create one, select *Create an Intermediate Certificate Authority* from the Method drop-down.

Note: The higher-level CA must already exist on AZTCO-FW® (Created or imported)

- Choose the higher-level CA to sign this CA using the Signing Certificate Authority drop-down. Only CAs with private keys present will be shown, as this is required to properly sign this new CA.
- Fill in the remaining parameters identical to those for *Create an Internal Certificate Authority*.

7.1.2 Edit a Certificate Authority

After a CA has been added, it can be edited from the list of CAs found at System > Cert Manager on the CAs tab.

To edit a CA, click the ⧉ icon at the end of its row. The screen presented allows editing the fields as if the CA were being imported.

For information on the fields on this screen, see *Import an Existing Certificate Authority*. In most cases the purpose of this screen would be to correct the Serial of the CA if needed, or to add a key to an imported CA so it can be used to create and sign certificates and CRLs.

7.1.3 Export a Certificate Authority

From the list of CAs at System > Cert Manager on the CAs tab, the certificate and/or private key for a CA can be exported. In most cases the private key for a CA would not be exported, unless the CA is being moved to a new location or a backup is being made. When using the CA for a VPN or most other purposes, only export the certificate for the CA.

> Warning: If the private key for a CA gets into the wrong hands, the other party could generate new certificates that would be considered valid against the CA.

12.1. Certificate Authority Management

To export the certificate for a CA, click the ✹ icon on the *left*. To export the private key for the CA, click the 🔑 icon on the *right*. Hover the mouse pointer over the icon and a tooltip will display the action to be performed for easy confirmation. The files will download with the descriptive name of the CA as the file name, with the extension *.crt* for the certificate, and *.key* for the private key.

7.1.4 Remove a Certificate Authority

To remove a CA, first it must be removed from active use.

- Check areas that can use a CA, such as OpenVPN, IPsec, and packages.
- Remove entries utilizing the CA or select a different CA.
- Navigate to System > Cert Manager on the CAs tab.
- Find the CA to delete in the list.
- Click 🗑 at the end of the row for the CA.
- Click OK on the confirmation dialog.

If an error appears, follow the on-screen instructions to correct the problem and then try again.

7.2 Certificate Management

Certificates are managed from System > Cert Manager, on the Certificates tab. From this screen Certificates may be added, edited, exported, or deleted.

7.2.1 Create a new Certificate

To create a new certificate, start the process as follows:

- Navigate to System > Cert Manager on the Certificates tab.
- Click Add to create a new certificate.
- Enter a Descriptive name for the certificate. This is used as a label for this certificate throughout the GUI.

- Select the Method that best suits how the certificate will be generated. These options and further instructions are in the corresponding sections below:
  - Import an Existing Certificate
  - Create an Internal Certificate
  - Create a Certificate Signing Request

Import an Existing Certificate

If an existing certificate from an external source needs to be imported, it can be done by selecting the Method of *Import an Existing Certificate*. This can be useful for certificates that have been made using another system or for certificates that have been provided by a third party.

- Enter the Certificate data, this is required. It is typically contained in a file ending with *.crt*. It would be plain text, and enclosed in a block such as:

```
-----BEGIN CERTIFICATE-----
[A bunch of random-looking base64-encoded data]
-----END CERTIFICATE-----
```

- Enter the Private key data which is also required. This is typically in a file ending in *.key*. It would be plain text data enclosed in a block such as:

```
-----BEGIN RSA PRIVATE KEY-----
[A bunch of random-looking base64-encoded data]
-----END RSA PRIVATE KEY-----
```

- Click Save to finish the import process.

If any errors are encountered, follow the on-screen instructions to resolve them. The most common error is not pasting in the right portion of the certificate or private key. Make sure to include the entire block, including the beginning header and ending footer around the encoded data.

Create an Internal Certificate

The most common Method is *Create an Internal Certificate*. This will make a new certificate using one of the existing Certificate Authorities.

- Select the Certificate Authority by which this certificate will be signed. Only a CA that has a private key present can be in this list, as the private key is required in order for the CA to sign a certificate.
- Select the Key length to choose how “strong” the certificate is in terms of encryption. The longer the key, the more secure it is. However, longer keys can take more CPU time to process, so it isn’t always wise to use the maximum value. The default value of *2048* is a good balance.
- Select a Digest Algorithm from the supplied list. The current best practice is to use an algorithm stronger than SHA1 where possible. *SHA256* is a good choice.

  Note: Some older or less sophisticated equipment, such as VPN-enabled VoIP handsets may only support SHA1 for the Digest Algorithm. Consult device documentation for specifics.

- Select a Certificate Type which matches the purpose of this certificate.

  - Choose Server Certificate if the certificate will be used in a VPN server or HTTPS server. This indicates inside the certificate that it may be used in a server role, and no other.

    Note: Server type certificates include Extended Key Usage attributes indicating they may be used for Server Authentication as well as the OID 1.3.6.1.5.5.8.2.2 which is used by Microsoft to signifiy that a certificate may be used as an IKE intermediate. These are required for Windows 7 and later to trust the server certificate for use with certain types of VPNs. They also are marked with a constraint indicating that they are not a CA, and have nsCertType set to "server".

  - Choose User Certificate if the certificate can be used in an end-user capacity, such as a VPN client, but it cannot be used as a server. This prevents a user from using their own certificate to impersonate a server.

    Note: User type certificates include Extended Key Usage attributes indicating they may be used for client authentication. They also are marked with a constraint indicating that they are not a CA.

  - Choose Certificate Authority to create an intermediate CA. A certificate generated in this way will be subordinate to the chosen CA. It can create its own certificates, but the root CA must also be included when it is used. This is also known as "chaining".

- Enter a value for Lifetime to specify the number of days for which the certificate will be valid. The duration depends on personal preferences and site policies. Changing the certificate frequently is more secure, but it is also a management headache as it requires reissuing new certificates when they expire. By default the GUI suggests using 3650 days, which is approximately 10 years.

- Enter values for the Distinguished name section for personalized parameters in the certificate. Most of these fields will be pre-populated with data from the CA. These are typically filled in with an organization's information, or in the case of an individual, personal information. This information is mostly cosmetic, and used to verify the accuracy of the certificate, and to distinguish one certificate from another. Punctuation and special characters must not be used.

  - Select the Country Code from the list. This is the ISO-recognized country code, not a hostname top-level domain.

  - Enter the State or Province fully spelled out, not abbreviated.

  - Enter the City.

  - Enter the Organization name, typically the company name.

  - Enter a valid Email Address.

  - Enter the Common Name (CN). This field is the internal name that identifies the certificate. Unlike a CA, the CN for a certificate should be a username or hostname. For instance, it could be called *VPNCert*, *user01*, or *vpnrouter.example.com*.

    Note: Although it is technically valid, avoid using spaces in the CN.

- Click + Add to add Alternative Names if they are required. Alternative Names allow the certificate to specify multiple names that are all valid for the CN, such as two different hostnames, an additional IP address, a URL, or an e-mail address. This field may be left blank if it is not required or its purpose is unclear.

  - Enter a Type for the Alternative Name. This must contain one of *DNS* (FQDN or Hostname), *IP* (IP address), *URI* , or *email* .
  - Enter a Value for the Alternative Name. This field must contain an appropriately formatted value based on the Type entered.
  - Click 🗑 Delete at the end of the row for an unneeded Alternative Name.
  - Repeat this process for each additional Alternative Name.
- Click Save.

If errors are reported, such as invalid characters or other input problems, they will be described on the screen. Correct the errors, and attempt to Save again.

Create a Certificate Signing Request

Choosing a Method of *Certificate Signing Request* creates a new request file that can be sent into a third party CA to be signed. This would be used to obtain a certificate from a trusted root certificate authority. Once this Method has been chosen, the remaining parameters for creating this certificate are identical to those for *Create an Internal Certificate*.

12.2.2 Export a Certificate

From the list of certificates at System > Cert Manager on the Certificates tab, a certificate and/or its private key may be exported.

To export the certificate, click the ✹ icon. To export the private key for the certificate, click the 🔑 icon. To export the CA certificate, certificate and the private key for the certificate together in a PKCS#12 🗃 file, click the icon. To confirm the proper file is being exported, hover the mouse pointer over the icon and a tooltip will display the action to be performed.

The files will download with the descriptive name of the certificate as the file name, and the extension *.crt* for the certificate and *.key* for the private key, or *.p12* for a PKCS#12 file.

7.2.3 Remove a Certificate

To remove a certificate, first it must be removed from active use.

- Check areas that can use a certificate, such as the WebGUI options, OpenVPN, IPsec, and packages.
- Remove entries using the certificate, or choose another certificate.
- Navigate to System > Cert Manager on the Certificates tab.
- Locate the certificate to delete in the list
- Click 🗑 at the end of the row for the certificate.
- Click OK on the confirmation dialog.

If an error appears, follow the on-screen instructions to correct the problem and then try again.

7.2.4 User Certificates

If a VPN is being used that requires user certificates, they may be created in one of several ways. The exact method depends on where the authentication for the VPN is being performed and whether or not the certificate already exists.

12.2. Certificate Management

No Authentication or External Authentication

If there is no user authentication, or if the user authentication is being performed on an external server (RADIUS, LDAP, etc) then make a user certificate like any other certificate described earlier. Ensure that *User Certificate* is selected for the Certificate Type and set the Common Name to be the user's username.

Local Authentication / Create Certificate When Creating a User

If user authentication is being performed on AZTCO-FW® software, the user certificate can be made inside of the User Manager.

- Navigate to System > User Manager
- Create a user. See *User Management and Authentication* for details.
- Fill in the Username and Password
- Select Click to create a user certificate in the User Certificates section, which will display a simple form for creating a user certificate.
  - Enter a short Descriptive Name, which can be the username or something such as *Bob's Remote Access VPN Cert*.
  - Choose the proper Certificate Authority for the VPN.
  - Adjust the Key Length and Lifetime if desired.
- Finish any other required user details.
- Click Save

Local Authentication / Add a Certificate to an Existing User

To add a certificate to an existing user:

- Navigate to System > User Manager

  ⧉
  - Clickto edit the user

  ✚
  - ClickAdd under User Certificates.
- Choose options as needed available from the certificate creation process described in *Create a new Certificate*, or select *Choose an existing certificate* and then select from the Existing Certificates

For more information on adding and managing users, see *User Management and Authentication*.

7.3 Certificate Revocation List Management

Certificate Revocation Lists (CRLs) are a part of the X.509 system that publish lists of certificates that should no longer be trusted. These certificates may have been compromised or otherwise need to be invalidated. An application using a CA, such as OpenVPN may optionally use a CRL so it can verify connecting client certificates. A CRL is generated and signed against a CA using its private key, so in order to create or add certificates to a CRL in the GUI, the private key of the CA must be present. If the CA is managed externally and the private key for the CA is not on the firewall, a CRL may still be generated outside of the firewall and imported.

The traditional way to use a CRL is to only have one CRL per CA and only add invalid certificates to that CRL. In AZTCO-FW®, however, multiple CRLs may be created for a single CA. In OpenVPN, different CRLs may be chosen for separate VPN instances. This could be used, for example, to prevent a specific certificate from connecting to one instance while allowing it to connect to another. For IPsec, all CRLs are consulted and there is no selection as currently exists with OpenVPN.

Certificate Revocation Lists are managed from System > Cert Manager, on the Certificate Revocation tab. From this screen CRL entries can be added, edited, exported, or deleted. The list will show all Certificate Authorities and an option to add a CRL. The screen also indicates whether the CRL is internal or external (imported), and it shows a count of how many certificates have been revoked on each CRL.

Note: CRLs generated using AZTCO-FW software version 2.2.4-RELEASE and later properly include the authorityKeyIdentifier attribute to allow proper functionality with strongSwan for use with IPsec.

7.3.1 Create a new Certificate Revocation List

To create a new CRL:

- Navigate to System > Cert Manager, on the Certificate Revocation tab.
- Find the row with the CA that the CRL will be created for.
- Click + Add or Import CRL at the end of the row to create a new CRL.
- Choose *Create an Internal Certificate Revocation List* for the Method.
- Enter a Descriptive Name for the CRL, which is used to identify this CRL in lists around the GUI. It's usually best to include a reference to the name of the CA and/or the purpose of the CRL.
- Select the proper CA from the Certificate Authority drop-down menu.
- Enter the number of days for which the CRL should be valid in the Lifetime box. The default value is *9999* days, or almost 27 and a half years.
- Click Save

The browser will be return to the CRL list, and the new entry will be shown there.

7.3.2 Import an Existing Certificate Revocation List

To import a CRL from an external source:

- Navigate to System > Cert Manager, on the Certificate Revocation tab
- Find the row with the CA that the CRL will be imported for.
- Click + Add or Import CRL at the end of the row to create a new CRL.

- Choose *Import an Existing Certificate Revocation List* for the Method.
- Enter a Descriptive Name for the CRL, which is used to identify this CRL in lists around the GUI. It's usually best to include a reference to the name of the CA and/or the purpose of the CRL.
- Select the proper CA from the Certificate Authority drop-down menu.
- Enter the CRL data. This is typically in a file ending in *.crl*. It would be plain text data enclosed in a block such as:

```
-----BEGIN X509 CRL-----
[A bunch of random-looking base64-encoded data]
-----END X509 CRL-----
```

- Click Save to finish the import process.

If an error appears, follow the on-screen instructions to correct the problem and then try again. The most common error is not pasting in the right portion of the CRL data. Make sure to enter the entire block, including the beginning header and ending footer around the encoded data.

7.3.3 Export a Certificate Revocation List

From the list of CRLs at System > Cert Manager on the Certificate Revocation tab, a CRL may also be exported.

To export the CRL, click the 📥 icon. The file will download with the descriptive name of the CRL as the file name, and the extension *.crl*.

7.3.4 Delete a Certificate Revocation List

- Check areas that can use a CRL, such as OpenVPN.
- Remove entries using the CRL, or choose another CRL instead.
- Navigate to System > Cert Manager on the Certificate Revocation tab.
- Locate the CRL to delete in the list
- Click the 🗑 icon at the end of the row for the CRL.
- Click OK on the confirmation dialog.

If an error appears, follow the on-screen instructions to correct the problem and then try again.

7.3.5 Revoke a Certificate

A CRL isn't very useful unless it contains revoked certificates. A certificate is revoked by adding the certificate to a CRL:

- Navigate to System > Cert Manager on the Certificate Revocation tab.
- Locate the CRL to edit in the list

- Click the ⧉ icon at the end of the row for the CRL. A screen will be presented that lists any currently revoked certificates, and a control to add new ones.
- Select the certificate from the Choose a Certificate to Revoke list.
- Select a Reason from the drop-down list to indicate why the certificate is being revoked. This information doesn't affect the validity of the certificate it is merely informational in nature. This option may be left at the default value.
- Click Add and the certificate will be added to the CRL.

Certificates can be removed from the CRL using this screen as well:

- Navigate to System > Cert Manager on the Certificate Revocation tab.
- Locate the CRL to edit in the list
- Click the ⧉ icon at the end of the row for the CRL.
- Find the certificate in the list and click the 🗑 icon to remove it from the CRL.
- Click OK on the confirmation dialog.

After adding or removing a certificate, the CRL will be re-written if it is currently in use by any VPN instances so that the CRL changes will be immediately active.

7.3.6 Updating an Imported Certificate Revocation List

To update an imported CRL:

- Navigate to System > Cert Manager on the Certificate Revocation tab.
- Locate the CRL to edit in the list
- Click the ⧉ icon at the end of the row for the CRL.
- Erase the pasted content in the CRL Data box and replace it with the contents of the new CRL
- Click Save.

After updating the imported CRL, it will be re-written if it is currently in use by any VPN instances so that the CRL changes will be immediately active.

7.4 DH Parameters

To put it simply, the DH parameters are extra bits of randomness that help out during the key exchange process. They do not have to match on both sides of the tunnel, and new DH parameters can be made at any time. DH parameters are not specific to a given setup in the way that certificates or keys are. There is no need to import an existing set of DH parameters because generating new parameters is a better practice.

AZTCO-FW® software ships with a default set of DH parameter files so that new firewalls do not have to spend significant CPU resources to build them when they are needed. These pre-generated parameters are stored in

/etc/dh-parameters.. Selecting a specific length in the GUI will use the DH parameter set from the corresponding file. These DH parameters are not stored in config.xml.

To generate a new set of DH parameters, which can take quite a long time depending on the hardware in use, run the following commands:

```
/usr/bin/openssl dhparam -out /etc/dh-parameters.1024 1024
/usr/bin/openssl dhparam -out /etc/dh-parameters.2048 2048
/usr/bin/openssl dhparam -out /etc/dh-parameters.4096 4096
```

CPU time used to generate the parameters increases significantly with length. For example, generating 1024-bit DH parameters only takes about 7 seconds on a C2758 CPU, but generating 2048-bit parameters takes 4 minutes, and generating 4096-bit parameters takes 10 minutes.

The AZTCO-FW webGUI will allow longer DH parameter to be selected if they exist in /etc/ in the format specified above.

Supported lengths are: 1024, 2048, 3072, 4096, 7680, 8192, 15360, and 16384

For example, to generate a new set of DH parameters of length 8192, run:

12.4. DH Parameters

```
/usr/bin/openssl dhparam -out /etc/dh-parameters.8192 8192
```

AZTCO-FW® software includes a central Certificate Manager under System > Cert Manager, used to create and maintain Certificate Authorities, Certificates, and Certificate Revocation Lists.

Entries in the Certificate Manager are used by the firewall for purposes such as TLS for the GUI, VPNs, LDAP, various packages, and more.

7.5 Basic Introduction to X.509 Public Key Infrastructure

One authentication option for VPNs is to use X.509 keys. An in depth discussion of X.509 and Public Key Infrastructure (PKI) is outside the scope of this documentation, and is the topic of a number of entire books for those interested in details. This chapter provides the very basic understanding necessary for creating and managing certificates in AZTCO-FW® software.

With PKI, first a Certificate Authority (CA) is created. This CA then signs all of the individual certificates in the PKI. The certificate of the CA is used on VPN servers and clients to verify the authenticity of server and client certificates used. The certificate for the CA can be used to verify signing on certificates, but not to sign certificates. Signing certificates requires the private key for the CA. The secrecy of the CA private key is what ensures the security of a PKI. Anyone with access to the CA private key can generate certificates to be used on a PKI, hence it must be kept secure. This key is never distributed to clients or servers.

Warning: Never copy more files to clients than are needed, as this may compromise the security of the PKI.

A certificate is considered valid if it has been trusted by a given CA. In this case of VPNs, this means that a certificate made from a specific CA would be considered valid for any VPN using that CA. For that reason the best practice is to create a unique CA for each VPN that has a different level of security. For instance, if there are two mobile access VPNs with the same security access, using the same CA for those VPNs is OK. However if one VPN is for users and another VPN is for remote management, each with different restrictions, then a unique CA for each VPN should be used.

Certificate Revocation Lists (CRLs) are lists of certificates that have been compromised or otherwise need to be invalidated. Revoking a certificate will cause it to be considered untrusted so long as the application using the CA also uses a CRL. CRLs are generated and signed against a CA using its private key, so in order to create or add certificates to a CRL in the GUI the private key for a CA must be present.

CHAPTER

EIGHT

FIREWALL

One of the primary functions performed by AZTCO-FW® software is filtering traffic. This chapter covers fundamentals of firewalling, best practices, and required information necessary to configure firewall rules.

8.1 Firewall

One of the primary purposes of AZTCO-FW® software is to act as a firewall, deciding which traffic to pass or block between networks.

8.1.1 Managing Firewall Rules

Firewall rules control traffic passing through the firewall. These topics describe how to create and manage rules, plus settings related to rules.

Firewalling Fundamentals

This section deals primarily with introductory firewall concepts and lays the ground work for understanding how to configure firewall rules using AZTCO-FW® software.

Basic Terminology

Rule and *ruleset* are two terms used throughout this chapter:

Rule Refers to a single entry on the Firewall > Rules screen. A rule instructs the firewall how to match or handle network traffic.

Ruleset Refers to a group of rules collectively. Either *all* firewall rules as a whole, or a set of rules in a specific context such as the rules on an interface tab. The complete firewall ruleset is the sum of all user configured and automatically added rules, which are covered further throughout this chapter.

Rulesets on the Interface tabs are evaluated on a first match basis by AZTCO-FW. This means that reading the ruleset for an interface from top to bottom, the first rule that matches will be the one used by the firewall. Evaluation stops after reaching this match and then the firewall takes the action specified by that rule. Always keep this in mind when creating new rules, especially when crafting rules to restrict traffic. The most permissive rules should be toward the bottom of the list, so that restrictions or exceptions can be made above them.

Note: The Floating tab is the lone exception to this rule processing logic. It is covered in a later section of this chapter.

Stateful Filtering

AZTCO-FW is a stateful firewall, which means it remembers information about connections flowing through the firewall so that reply traffic can be allowed automatically. This data is retained in the State Table. The connection information in the state table includes the source, destination, protocol, ports, and more: Enough to uniquely identify a specific connection.

Using this mechanism, traffic need only be permitted on the interface where it enters the firewall. When a connection matches a pass rule the firewall creates an entry in the state table. Reply traffic to connections is automatically allowed back through the firewall by matching it against the state table rather than having to check it against rules in both directions. This includes any related traffic using a different protocol, such as ICMP control messages that may be provided in response to a TCP, UDP, or other connection.

See also:

See *Firewall Advanced* and *State Type* for more information about state options and types.

State table size

The firewall state table has a maximum size to prevent memory exhaustion. Each state takes approximately 1 KB of RAM. The default state table size in AZTCO-FW is calculated by taking about 10% of the RAM available in the firewall by default. On a firewall with 1GB of RAM, the default state table size can hold approximately 100,000 entries.

See also:

See *Large State Tables* for more information on state table sizing and RAM usage.

Each user connection typically consists of two states: One created as it enters the firewall, and one as it leaves the firewall. Therefore, with a state table size of 1,000,000, the firewall can handle approximately 500,000 user sessions actively traversing the firewall before any additional connections will be dropped. This limit can be increased as needed so long as it does not exceed the available amount of RAM in the firewall.

To increase the state table size:

- Navigate to System > Advanced on the Firewall & NAT tab
- Enter the desired number for Firewall Maximum States, or leave the box blank for the default calculated value. See Figure *Increased State Table Size to 2,000,000*
- Click Save

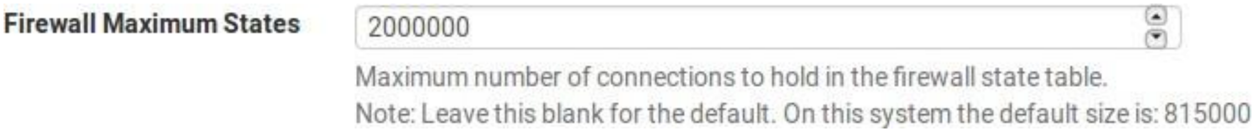

Firewall Maximum States 2000000

Maximum number of connections to hold in the firewall state table.
Note: Leave this blank for the default. On this system the default size is: 815000

Fig. 1: Increased State Table Size to 2,000,000

Historical state table usage is tracked by the firewall. To view the graph:

- Navigate to Status > Monitoring

- Click 🔧 to expand the graph options
- Set Category for the Left Axis to *System*
- Set the Graph for the Left Axis to *States*
- Click ↻ Update Graphs

Block vs. Reject

There are two ways to disallow traffic using firewall rules on AZTCO-FW: Block and reject.

A rule set to block will silently drop traffic. A blocked client will not receive any response and thus will wait until its connection attempt times out. This is the behavior of the default deny rule in AZTCO-FW.

A rule set to reject will respond back to the client for denied TCP and UDP traffic, letting the sender know that the connection was refused. Rejected TCP traffic receives a TCP RST (reset) in response, and rejected UDP traffic receives an ICMP unreachable message in response. Though reject is a valid choice for any firewall rule, IP protocols other than TCP and UDP are not capable of being rejected; These rules will silently drop other IP protocols because there is no standard for rejecting other protocols.

Deciding Between Block and Reject

There has been much debate amongst security professionals over the years as to the value of block vs. reject. Some argue that using block makes more sense, claiming it "slows down" attackers scanning the Internet. When a rule is set to reject, a response is sent back immediately that the port is closed, while block silently drops the traffic, causing the attacker's port scanner to wait for a response. That argument does not hold water because every good port scanner can scan hundreds or thousands of hosts simultaneously, and the scanner is not stalled waiting for a response from closed ports. There is a minimal difference in resource consumption and scanning speed, but so slight that it shouldn't be a consideration.

If the firewall blocks all traffic from the Internet, there is a notable difference between block and reject: Nobody knows the firewall is online. If even a single port is open, the value of that ability is minimal because the attacker can easily determine that the host is online and will also know what ports are open whether or not the blocked connections have been rejected by the firewall. While there isn't significant value in block over reject, we still recommend using block on WAN rules. There is some value in not actively handing information to potential attackers, and it is also a bad practice to automatically respond to an external request unnecessarily.

For rules on internal interfaces we recommend using reject in most situations. When a host tries to access a resource that is not permitted by firewall rules, the application accessing it may hang until the connection times out or the client program stops trying to access the service. With reject the connection is immediately refused and the client avoids these hangs. This is usually nothing more than an annoyance, but we still generally recommend using reject to avoid potential application problems induced by silently dropping traffic inside a network.

Introduction to the Firewall Rules screen

This section provides an introduction and overview of the Firewall Rules screen located at Firewall > Rules. This page lists the WAN ruleset to start with, which by default has no entries other than those for Block private networks and Block bogon networks if those options are active on the WAN interface, as shown in Figure *Default WAN Rules*.

Tip: Click ⚙ the to the right of the Block private networks or Block bogon networks rules to reach the WAN interface configuration page where these options can be enabled or disabled. (See *Block Private Networks* and *Block Bogon Networks* for more details.)

Click the LAN tab to view the LAN rules. By default, the only entries are the *Default allow LAN to any* rules for IPv4 and IPv6 as seen in Figure *Default LAN Rules*, and the Anti-Lockout Rule if it is active. The anti-lockout rule is designed to prevent administrators from accidentally locking themselves out of the GUI. Click ⚙ next to the anti-lockout rule to reach the page where this rule can be disabled.

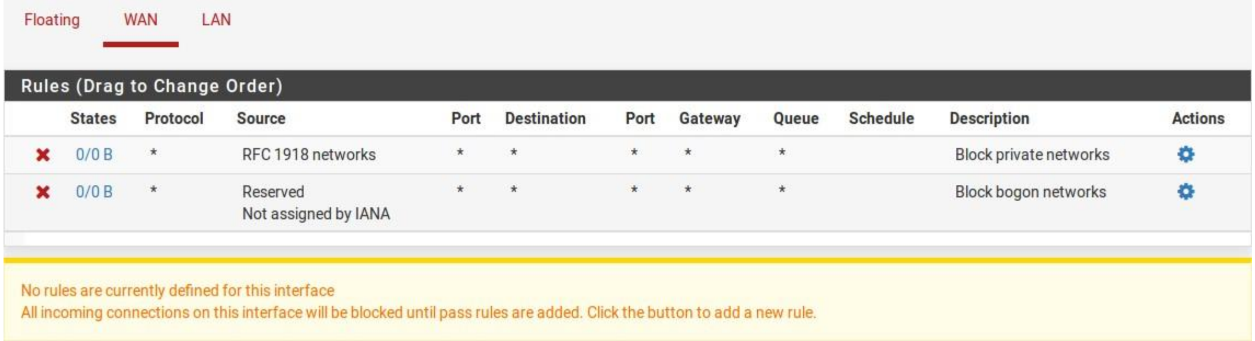

Rules (Drag to Change Order)

	States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
✖	0/0 B	*	RFC 1918 networks	*	*	*	*	*		Block private networks	⚙
✖	0/0 B	*	Reserved Not assigned by IANA	*	*	*	*	*		Block bogon networks	⚙

No rules are currently defined for this interface
All incoming connections on this interface will be blocked until pass rules are added. Click the button to add a new rule.

Fig. 2: Default WAN Rules

See also:

For more information on how the Anti-Lockout Rule works and how to disable the rule, see *Anti-lockout Rule* and *Anti-lockout*.

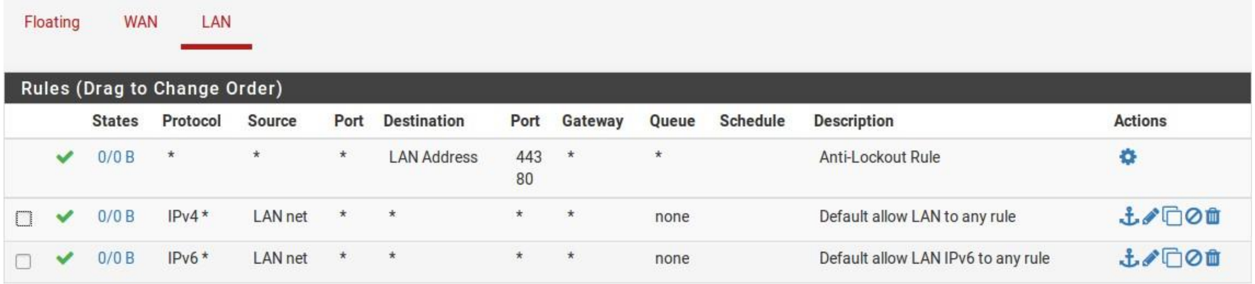

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
	✓	0/0 B	*	*	*	LAN Address	443 80	*	*		Anti-Lockout Rule	⚙
□	✓	0/0 B	IPv4 *	LAN net	*	*	*	*	none		Default allow LAN to any rule	⚓✎❐⊘🗑
□	✓	0/0 B	IPv6 *	LAN net	*	*	*	*	none		Default allow LAN IPv6 to any rule	⚓✎❐⊘🗑

Fig. 3: Default LAN Rules

To display rules for other interfaces, click their respective tabs. OPT interfaces will appear with their descriptive names, so if the OPT1 interface was renamed DMZ, then the tab for its rules will also say DMZ.

To the left of each rule is an indicator icon showing the action of the rule: pass (✖), block (✖), or reject ✋ (). If logging is enabled for the rule, ☰ is shown in the same area. If the rule has any advanced options enabled, an ⚙ icon is also displayed. Hovering the mouse cursor over any of these icons will display text explaining their meaning. The same icons are shown for disabled rules, except the icon and the rule are a lighter shade of their original color.

Adding a firewall rule

To add a rule to the top of the list, click ↴ Add.

To add a rule to the bottom of the list, click ↴ Add.

To make a new rule that is similar to an existing rule, click ⧉ to the right of the existing rule. The edit screen will appear with the existing rule's settings pre-filled, ready to be adjusted. When duplicating an existing rule, the new rule will be added directly *below* the original rule. For more information about how to configure the new rule, see *Configuring firewall rules*.

Editing Firewall Rules

To edit a firewall rule, click ⧉ to the right of the rule, or double click anywhere on the line.

The edit page for that rule will load, and from there adjustments are possible. See *Configuring firewall rules* for more information on the options available when editing a rule.

Moving Firewall Rules

Rules may be reordered in two different ways: Drag-and-drop, and using select-and-click.

To move rules using the drag-and-drop method:

- Move the mouse over the firewall rule to move, the cursor will change to indicate movement is possible.
- Click and hold the mouse button down
- Drag the mouse to the desired location for the rule
- Release the mouse button
- Click 💾 Save to store the new rule order

> Warning: Attempting to navigate away from the page after moving a rule, but before saving the rule, will result in the browser presenting an error confirming whether or not to exit the page. If the browser navigates away from the page without saving, the rule will still be in its original location.

To move rules in the list in groups or by selecting them first, use the select-and-click method:

- Check the box next to the left of the rules which need to be moved, or single click the rule. When the rule is selected, it will change color.
- Click ⚓ on the row below where the rule should be moved.

> **Tip:** Hold Shift before clicking the mouse on ⚓ to move the rule below the selected rule instead of above.

When moving rules using the select-and-click method, the new order is stored automatically.

Deleting Firewall Rules

To delete a single rule, click 🗑 to the right of the rule. The firewall will present a confirmation prompt before deleting the rule.

To delete multiple rules, check the box at the start of the rows that should be removed, then click the 🗑 Delete button at the bottom of the list. Rules may also be selected by single clicking anywhere on their line.

Disabling and Enabling Firewall Rules

To disable a rule, ☑ click at the end of its row. The appearance of the rule will change to a lighter shade to indicate ☑ that it is disabled and theicon changes to 🗑.

To enable a rule which was previously disabled, click 🗑 at the end of its row. The appearance of the rule will return to normal and the enable/disable icon will return to the original ☑.

A rule may also be disabled or enabled by editing the rule and toggling the Disabled checkbox.

Rule Separators

Firewall Rule Separators are colored bars in the ruleset that contain a small bit of text, but do not take any action on traffic. They are useful for visually separating or adding notes to special parts of the ruleset. Figure *Firewall Rule Separators Example* shows how they can be utilize to group and document the ruleset.

To create a new Rule Separator:

- Open the firewall rule tab where the Rule Separator will reside
- Click ➕ Separator

- Enter description text for the Rule Separator
- Choose the color for the Rule Separator by clicking the ● icon of the desired color
- Click and drag the Rule Separator to its new location
- Click 💾 Save inside the Rule Separator to store its contents

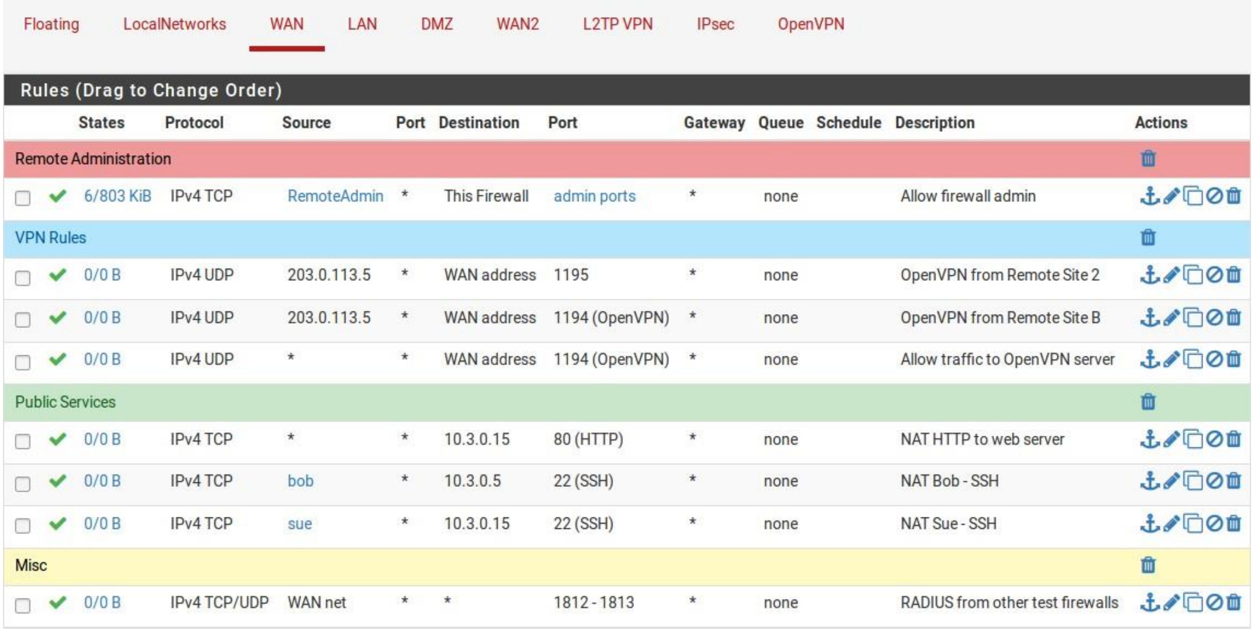

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
Remote Administration												🗑
☐	✓	6/803 KiB	IPv4 TCP	RemoteAdmin	*	This Firewall	admin ports	*	none		Allow firewall admin	⚓✎⧉⊘🗑
VPN Rules												🗑
☐	✓	0/0 B	IPv4 UDP	203.0.113.5	*	WAN address	1195	*	none		OpenVPN from Remote Site 2	⚓✎⧉⊘🗑
☐	✓	0/0 B	IPv4 UDP	203.0.113.5	*	WAN address	1194 (OpenVPN)	*	none		OpenVPN from Remote Site B	⚓✎⧉⊘🗑
☐	✓	0/0 B	IPv4 UDP	*	*	WAN address	1194 (OpenVPN)	*	none		Allow traffic to OpenVPN server	⚓✎⧉⊘🗑
Public Services												🗑
☐	✓	0/0 B	IPv4 TCP	*	*	10.3.0.15	80 (HTTP)	*	none		NAT HTTP to web server	⚓✎⧉⊘🗑
☐	✓	0/0 B	IPv4 TCP	bob	*	10.3.0.5	22 (SSH)	*	none		NAT Bob - SSH	⚓✎⧉⊘🗑
☐	✓	0/0 B	IPv4 TCP	sue	*	10.3.0.15	22 (SSH)	*	none		NAT Sue - SSH	⚓✎⧉⊘🗑
Misc												🗑
☐	✓	0/0 B	IPv4 TCP/UDP	WAN net	*	*	1812 - 1813	*	none		RADIUS from other test firewalls	⚓✎⧉⊘🗑

Fig. 4: Firewall Rule Separators Example

- Click 💾 Save at the bottom of the rule list

To move a Rule Separator:

- Open the firewall rule tab containing the Rule Separator
- Click and drag the Rule Separator to its new location
- Click 💾 Save at the bottom of the rule list

To delete a Rule Separator:

- Open the firewall rule tab containing the Rule Separator
- Click 🗑 inside the Rule Separator on the right side
- Click 💾 Save at the bottom of the rule list

Rule Separators cannot be edited. If a change in text or color is required, create a new Rule Separator and delete the existing entry.

Tracking Firewall Rule Changes

When a rule is created or updated the firewall records the user's login name, IP address, and a timestamp on the rule to track who added and/or last changed the rule in question. If the firewall automatically created the rule, that is also noted. This is done for firewall rules as well as port forwards and outbound NAT rules. An example of a rule update tracking block is shown in Figure *Firewall Rule Time Stamps*, which is visible when editing a firewall rule at the very bottom of the rule editing screen.

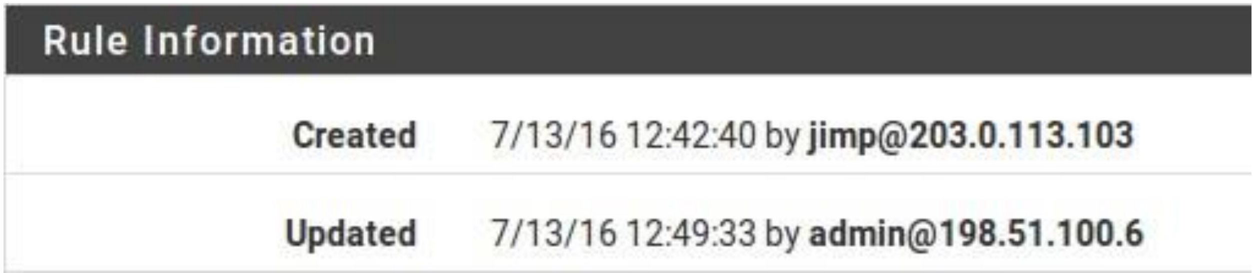

Rule Information

Created	7/13/16 12:42:40 by **jimp@203.0.113.103**
Updated	7/13/16 12:49:33 by **admin@198.51.100.6**

Fig. 5: Firewall Rule Time Stamps

Ingress Filtering

Ingress filtering refers to the concept of firewalling traffic entering a network from an external source such as the Internet. In deployments with multi-WAN, the firewall has multiple ingress points. The default ingress policy on AZTCO-FW® software is to block all traffic as there are no allow rules on WAN in the default ruleset. Replies to traffic initiated from inside the local network are automatically allowed to return through the firewall by the state table.

Egress Filtering

Egress filtering refers to the concept of firewalling traffic initiated inside the local network, destined for a remote network such as the Internet. AZTCO-FW, like nearly all similar commercial and open source solutions, comes with a LAN rule allowing everything from the LAN out to the Internet. This isn't the best way to operate, however. It has become the de facto default in most firewall solutions because it is what most people expect. The common misperception is "Anything on the internal network is 'trustworthy', so why bother filtering"?

Why employ egress filtering?

From our experience in working with countless firewalls from numerous vendors across many different organizations, most small companies and home networks do not employ egress filtering. It can increase the administrative burden as each new application or service may require opening additional ports or protocols in the firewall. In some environments it is difficult because the administrators do not completely know what is happening on the network, and they are hesitant to break things. In other environments it is impossible for reasons of workplace politics. The best practice is for administrators to configure the firewall to allow only the minimum required traffic to leave a network where possible. Tight egress filtering is important for several reasons:

Limit the Impact of a Compromised System

Egress filtering limits the impact of a compromised system. Malware commonly uses ports and protocols that are not required on most business networks. Some bots rely on IRC connections to phone home and receive instructions. Some will use more common ports such as TCP port 80 (normally HTTP) to evade egress filtering, but many do not.

If access to TCP port 6667, the usual IRC port, is not permitted by the firewall, bots that rely on IRC to function may be crippled by the filtering.

Another example is a case we were involved in where the inside interface of a AZTCO-FW installation was seeing 50-60 Mbps of traffic while the WAN had less than 1 Mbps of throughput. There were no other interfaces on the firewall. Some investigation showed the cause as a compromised system on the LAN running a bot participating in a distributed denial of service (DDoS) attack against a Chinese gambling web site. The attack used UDP port 80, and in this network UDP port 80 was not permitted by the egress ruleset so all the DDoS was accomplishing was stressing the inside interface of the firewall with traffic that was being dropped. In this situation, the firewall was happily chugging along with no performance degradation and the network's administrator did not know it was happening until it was discovered by accident.

The attack described in the above paragraph likely used UDP port 80 for two main reasons:

- UDP allows large packets to be sent by the client without completing a TCP handshake. With stateful firewalls being the norm, large TCP packets will not pass until the handshake is successfully completed, and this limits the effectiveness of the DDoS.
- Those who do employ egress filtering are commonly too permissive, allowing TCP and UDP where only TCP is required, as in the case of HTTP.

These types of attacks are commonly launched from compromised web servers. With a wide open egress ruleset, the traffic will go out to the Internet, and has the potential to overflow the state table on the firewall, cost money in bandwidth usage, and/or degrade performance for everything on the Internet connection.

Outbound SMTP is another example. Only allow SMTP (TCP port 25) to leave any network from a mail server. Or if a mail server is externally hosted, only allow internal systems to talk to that specific outside system on TCP port 25. This prevents every other system in the local network from being used as a spam bot, since their SMTP traffic will be dropped. Many mail providers have moved to using only authentication submission from clients using TCP port 587, so clients should not need access to port 25. This has the obvious benefit of limiting spam, and also prevents the network from being added to numerous black lists across the Internet that will prevent that site from sending legitimate e-mail to many mail servers. This may also prevent the ISP for that site from shutting off its Internet connection due to abuse.

The ideal solution is to prevent these types of things from happening in the first place, but egress filtering provides another layer that can help limit the impact if other measures fail.

Prevent a Compromise

Egress filtering can prevent a compromise in some circumstances. Some exploits and worms require outbound access to succeed. An older but good example of this is the Code Red worm from 2001. The exploit caused affected systems to pull an executable file via TFTP (Trivial File Transfer Protocol) and then execute it. A web server almost certainly does not need to use the TFTP protocol, and blocking TFTP via egress filtering prevented infection with Code Red even on unpatched servers. This is largely only useful for stopping completely automated attacks and worms as a real human attacker will find any holes that exist in egress filtering and use them to their advantage. Again, the correct solution to prevent such a compromise is to fix the network vulnerabilities used as an attack vector, however egress filtering can help.

Limit Unauthorized Application Usage

Many applications such as VPN clients, peer-to-peer software, instant messengers, and more rely on atypical ports or protocols to function. While a growing number of peer-to-peer and instant messenger applications will port hop until finding a port which is allowed out of the local network, many will be prevented from functioning by a restrictive egress ruleset, and this is an effective means of limiting many types of VPN connectivity.

Prevent IP Spoofing

This is a commonly cited reason for employing egress filtering, but AZTCO-FW automatically blocks spoofed traffic via pf's *antispoof* functionality, so it isn't applicable here. Preventing IP Spoofing means that malicious clients cannot send traffic with obviously falsified source addresses.

Prevent Information Leaks

Certain protocols should never be allowed to leave a local network. Specific examples of such protocols vary from one environment to another, but a few common examples are:

- Microsoft RPC (Remote Procedure Call) on TCP port 135
- NetBIOS on TCP and UDP ports 137 through 139
- SMB/CIFS (Server Message Block/Common Internet File System) on TCP and UDP port 445.

Stopping these protocols can prevent information about the internal network from leaking onto the Internet, and will prevent local systems from initiating authentication attempts with Internet hosts. These protocols also fall under *Limit the Impact of a Compromised System* as discussed previously since many worms have relied upon these protocols to function. Other protocols that may be relevant are syslog, SNMP, and SNMP traps. Restricting this traffic will prevent misconfigured network devices from sending logging and other potentially sensitive information out to the Internet. Rather than worry about what protocols can leak information out of a local network and need to be blocked, the best practice is to only allow the traffic that is required.

Approaches for implementing egress filtering

On a network that has historically not employed egress filtering, it can be difficult to know what traffic is absolutely necessary. This section describes some approaches for identifying traffic and implementing egress filtering.

Allow what is known, block the rest, and work through the fallout

One approach is to add firewall rules for known required traffic to be permitted. Start with making a list of things known to be required such as in Table *Egress Traffic Required*.

Table 1: Egress Traffic Required

Description	Source	Destination	Destination port
HTTP and HTTPS from all hosts	LAN Network	Any	TCP 80 and 443
SMTP from mail server	Mail Server	Any	TCP 25
DNS queries from internal DNS servers	DNS Servers	Any	TCP and UDP 53

After making the list, configure firewall rules to pass only that traffic and let everything else hit the default deny rule.

Log Traffic and Analyze Logs

Another alternative is to enable logging on all pass rules and send the logs to a syslog server. The logs can be analyzed by the syslog server to see what traffic is leaving the network. AZTCO-FW uses a custom log format, so the logs typically need be parsed by a custom script unless the server has some knowledge of the AZTCO-FW filter log format. Analysis of the logs will help build the required ruleset with less fallout as it will yield a better idea of what traffic is necessary on the local network.

Firewall Rule Best Practices

This section covers general best practices for firewall rule configuration.

Default Deny

There are two basic philosophies in computer security related to access control: default allow and default deny. A default deny strategy for firewall rules is the best practice. Firewall administrators should configure rules to permit only the bare minimum required traffic for the needs of a network, and let the remaining traffic drop with the default deny rule built into AZTCO-FW® software. In following this methodology, the number of deny rules in a ruleset will be minimal. They still have a place for some uses, but will be minimized in most environments by following a default deny strategy.

In a default two-interface LAN and WAN configuration, AZTCO-FW utilizes default deny on the WAN and default allow on the LAN. Everything inbound from the Internet is denied, and everything out to the Internet from the LAN is permitted. All home grade routers use this methodology, as do all similar open source projects and most similar commercial offerings. It's what most people expect out of the box, therefore it is the default configuration. That said, while it is a convenient way to start, it is not the recommended means of long-term operation.

AZTCO-FW users often ask "What bad things should I block?" but that is the wrong question as it applies to a default allow methodology. Noted security professional Marcus Ranum includes default permit in his "Six Dumbest Ideas in Computer Security" paper, which is recommended reading for any security professional. Permit only what a network requires and avoid leaving the default allow all rule on the LAN and adding block rules for "bad things" above the permit rule.

Keep it short

The shorter a ruleset, the easier it is to manage. Long rulesets are difficult to work with, increase the chances of human error, tend to become overly permissive, and are significantly more difficult to audit. Utilize aliases to keep the ruleset as short as possible.

Review Firewall Rules

We recommend a manual review of the firewall rules and NAT configuration on a periodic basis to ensure they still match the minimum requirements of the current network environment. The recommended frequency of such reviews varies from one environment to another. In networks that do not change frequently, with a small number of firewall administrators and good change control procedures, quarterly or semi-annually is usually adequate. For fast changing environments or those with poor change control and several people with firewall access, review the configuration at least on a monthly basis.

Quite often when reviewing rules with customers we ask about specific rules and they respond with "We removed that server six months ago." If something else would have taken over the same internal IP address as the previous server, then traffic would have been allowed to the new server that may not have been intended.

Document The Configuration

In all but the smallest networks, it can be hard to recall what is configured where and why. We always recommend using the Description field in firewall and NAT rules to document the purpose of the rules. In larger or more complex deployments, create and maintain a more detailed configuration document describing the entire AZTCO-FW configuration. When reviewing the firewall configuration in the future, this will help determine which rules are necessary and why they are there. This also applies to any other area of the configuration.

It is also important to keep this document up to date. When performing periodic configuration reviews, also review this document to ensure it remains up-to-date with the current configuration. Ensure this document is updated whenever configuration changes are made.

Reducing Log Noise

By default, AZTCO-FW will log packets blocked by the default deny rule. This means all of the noise getting blocked from the Internet will be logged. Sometimes there will not be much noise in the logs, but in many environments there will inevitably be something incessantly spamming the logs.

On networks using large broadcast domains – a practice commonly employed by cable ISPs – this is most often NetBIOS broadcasts from clue-deficient individuals who connect Windows machines directly to their broadband connections. These machines will constantly pump out broadcast requests for network browsing, among other things. ISP routing protocol packets may also be visible, or router redundancy protocols such as VRRP or HSRP. In co-location environments such as data centers, a combination of all of those things may be present.

Because there is no value in knowing that the firewall blocked 14 million NetBIOS broadcasts in the past day, and that noise could be covering up logs that are important, it is a good idea to add a block rule on the WAN interface for repeated noise traffic. By adding a block rule *without logging enabled* on the WAN interface, this traffic will still be blocked, but no longer fill the logs.

The rule shown in Figure *Firewall Rule to Prevent Logging Broadcasts* is configured on a test system where the "WAN" is on an internal LAN behind an edge firewall. To get rid of the log noise to see the things of interest, we added this rule to block – but not log – anything with the destination of the broadcast address of that subnet.

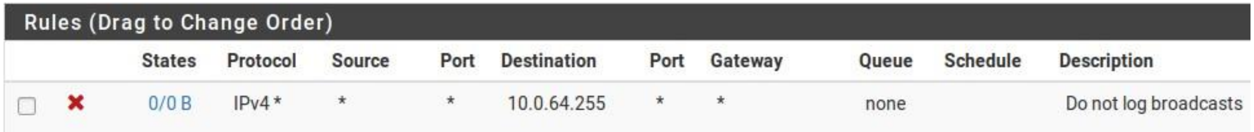

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
□	✖	0/0 B	IPv4 *	*	*	10.0.64.255	*	*	none		Do not log broadcasts

Fig. 6: Firewall Rule to Prevent Logging Broadcasts

We recommend adding similar rules, matching the specifics of any log noise observed in an environment. Check the firewall logs under Status > System Logs, Firewall tab to see what kind of traffic the firewall is blocking, and review how often it appears in the log. If any particular traffic is consistently being logged more than 5 times a minute, and the traffic is not malicious or noteworthy, add a block rule for it to reduce log noise.

Logging Practices

Out of the box, AZTCO-FW does not log any passed traffic and logs all dropped traffic. This is the typical default behavior of almost every open source and commercial firewall. It is the most practical, as logging all passed traffic is rarely desirable due to the load and log levels generated. This methodology is a bit backwards, however, from a security perspective. Blocked traffic cannot harm a network so its log value is limited, while traffic that gets passed could be very important log information to have if a system is compromised. After eliminating any useless block noise as described in the previous section, the remainder is of some value for trend analysis purposes. If significantly more or less log volume than usual is observed, it is probably good to investigate the nature of the logged traffic. OSSEC, an open source host-based intrusion detection system (IDS), is one system that can gather logs from AZTCO-FW via syslog and alert based on log volume abnormalities.

Rule Methodology

In AZTCO-FW® software, rules on interface tabs are applied on a per-interface basis, always in the inbound direction on that interface. This means traffic initiated from the LAN is filtered using the LAN interface rules. Traffic initiated from the Internet is filtered with the WAN interface rules. Because all rules in AZTCO-FW are stateful by default, a state table entry is created when traffic matches an allow rule. All reply traffic is automatically permitted by this state table entry.

The exception to this is Floating rules (*Floating Rules*), which can act on any interface using the inbound, outbound, or both directions. Outbound rules are never required, because filtering is applied on the inbound direction of every interface. In some limited circumstances, such as a firewall with numerous internal interfaces, having them available can significantly reduce the number of required firewall rules. In such a case, apply egress rules for Internet traffic as outbound rules on the WAN to avoid having to duplicate them for every internal interface. The use of inbound and outbound filtering makes a configuration more complex and more prone to user error, but it can be desirable in specific applications.

Interface Groups

Interface groups, discussed in *Interface Groups*, are a method to place rules on multiple interfaces at the same time. This can simplify some rule configurations if similar rules are required on many interfaces in the same way. Interface group rules, like interface rules, are processed in the inbound direction only. The VPN tabs for OpenVPN, L2TP, and the PPPoE server are all special Interface groups that are automatically created behind the scenes.

For example, a group may be used for a collection of interfaces including all LAN or DMZ type interfaces, or for a group of VLANs.

Note: Interface groups are not effective with Multi-WAN because group rules cannot properly handle reply-to. Due to that deficiency, traffic matching a group rule on a WAN that does not have the default gateway will go back out the WAN with the default gateway, and not through the interface which it entered.

Rule Processing Order

So far we have talked about how the rules are processed on an interface tab, but there are three main classes of rules: Regular interface rules, Floating rules, and Interface Group rules (including VPN tab rules). The order of processing of these types is significant, and it works like so:

1. Floating Rules
2. Interface Group Rules
3. Interface Rules

The rules are ordered in that way in the actual ruleset, keep that in mind when crafting rules. For example, if an interface group contains a rule to block traffic, that rule cannot be overridden with an interface tab rule because the traffic has already been acted upon by the group rule, which was matched first in the ruleset.

The rules are processed until a match is found, however, so if a packet is *not* matched in the group rules, it can still be matched by an interface rule.

Another significant place this comes into play is with assigned OpenVPN interfaces. If an “allow all” rule is in place on the OpenVPN tab, it is matched with the group rules. This means the rules on the interface tab will not apply. This can be a problem if OpenVPN rules need to have reply-to in order to ensure certain traffic exits back via the VPN. See also:

See *Ordering of NAT and Firewall Processing* for a more detailed analysis of rule processing and flow through the firewall, including how NAT rules come into play.

Automatically Added Firewall Rules

AZTCO-FW automatically adds internal firewall rules for a variety of reasons. This section describes automatically added rules and their purpose.

Anti-lockout Rule

To prevent locking an administrator out of the web interface, AZTCO-FW enables an anti-lockout rule by default. This is configurable on the System > Advanced page under Anti-lockout. This automatically added rule allows traffic from any source inside the network containing the rule, to any firewall administration protocol listening on the LAN IP address. For example, it grants access to TCP port 443 for the WebGUI, TCP port 80 for the GUI redirect, and TCP port 22 if SSH is enabled. If the WebGUI port has been changed, the configured port is the one allowed by the anti-lockout rule.

In security-conscious environments, the best practice is to disable this rule and configure the LAN rules so only an alias of trusted hosts can access the administrative interfaces of the firewall. A better practice yet is to not allow access from the LAN but only from an isolated administrative management network.

Restricting access to the administrative interface from LAN

First, to configure the firewall rules as desired to restrict access to the required management interface(s). In this typical use case example, both SSH and HTTPS are used for management, so create a ManagementPorts alias containing these ports (Figure *Alias for Management Ports*).

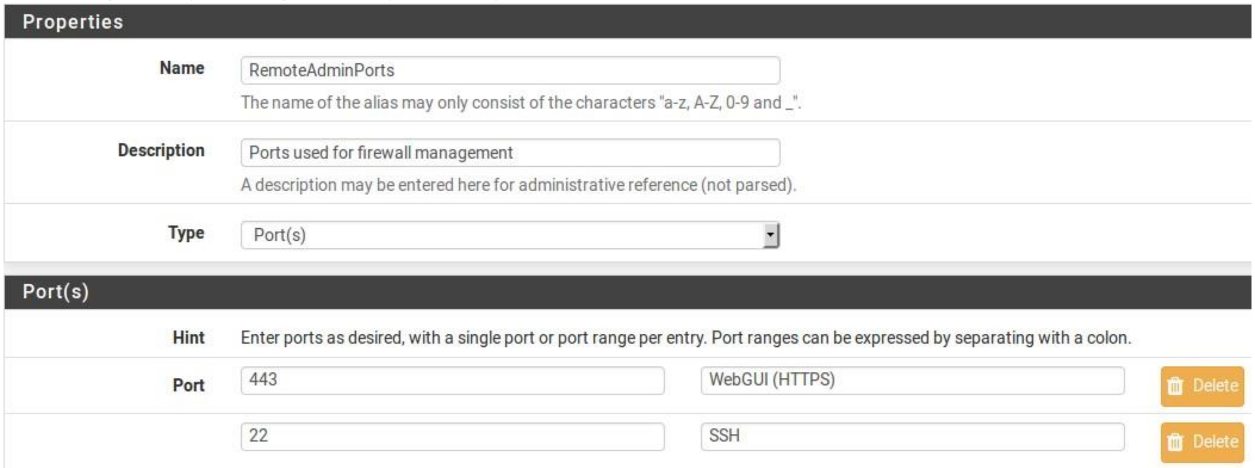

Properties

Name: RemoteAdminPorts
The name of the alias may only consist of the characters "a-z, A-Z, 0-9 and _".

Description: Ports used for firewall management
A description may be entered here for administrative reference (not parsed).

Type: Port(s)

Port(s)

Hint: Enter ports as desired, with a single port or port range per entry. Port ranges can be expressed by separating with a colon.

Port: 443 — WebGUI (HTTPS)

22 — SSH

Fig. 7: Alias for Management Ports

Then create an alias for hosts and/or networks that will have access to the management interfaces (Figure *Alias For Management Hosts*).

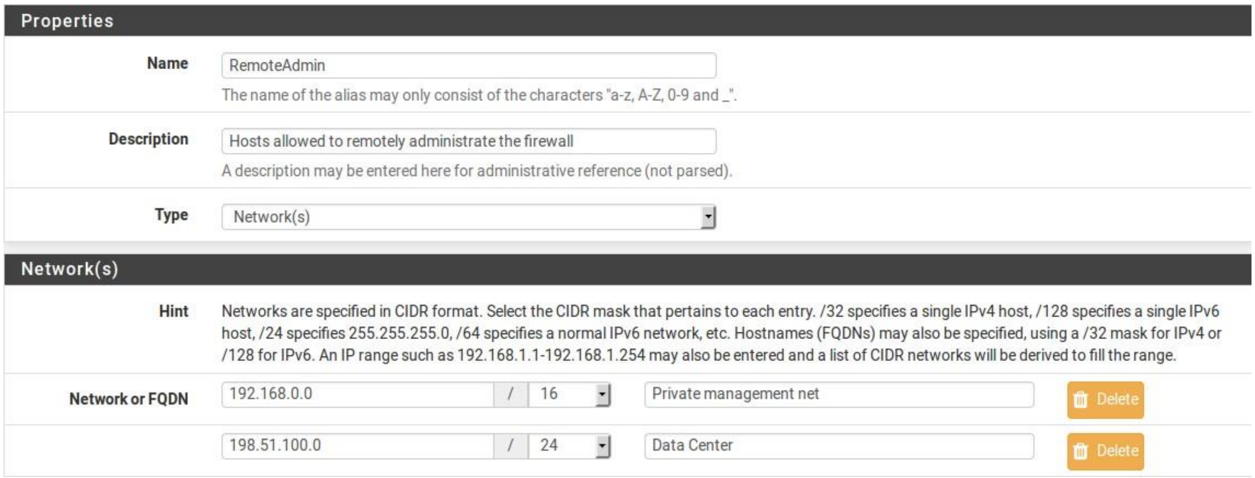

Properties

Name RemoteAdmin

The name of the alias may only consist of the characters "a-z, A-Z, 0-9 and _".

Description Hosts allowed to remotely administrate the firewall

A description may be entered here for administrative reference (not parsed).

Type Network(s)

Network(s)

Hint Networks are specified in CIDR format. Select the CIDR mask that pertains to each entry. /32 specifies a single IPv4 host, /128 specifies a single IPv6 host, /24 specifies 255.255.255.0, /64 specifies a normal IPv6 network, etc. Hostnames (FQDNs) may also be specified, using a /32 mask for IPv4 or /128 for IPv6. An IP range such as 192.168.1.1-192.168.1.254 may also be entered and a list of CIDR networks will be derived to fill the range.

Network or FQDN	192.168.0.0	/ 16	Private management net	Delete
	198.51.100.0	/ 24	Data Center	Delete

Fig. 8: Alias For Management Hosts

The resulting aliases are shown in Figure *Alias List*.

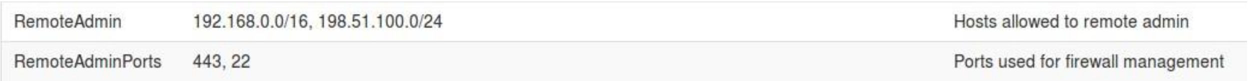

RemoteAdmin	192.168.0.0/16, 198.51.100.0/24	Hosts allowed to remote admin
RemoteAdminPorts	443, 22	Ports used for firewall management

Fig. 9: Alias List

Then the LAN firewall rules must be configured to allow access by the previously defined hosts, and deny access to all else. There are numerous ways to accomplish this, depending on specifics of the environment and how egress filtering is handled. Figure *Example Restricted Management LAN Rules* show two examples. The first allows DNS queries to the LAN IP address, which is needed if the DNS Resolver or DNS Forwarder are enabled, and also allows LAN hosts to ping the LAN IP address. It then rejects all other traffic. The second example allows access from the management hosts to the management ports, then rejects all other traffic to the management ports. Choose the methodology that works best for the network environment in question. Remember that the source port is not the same as the destination port.

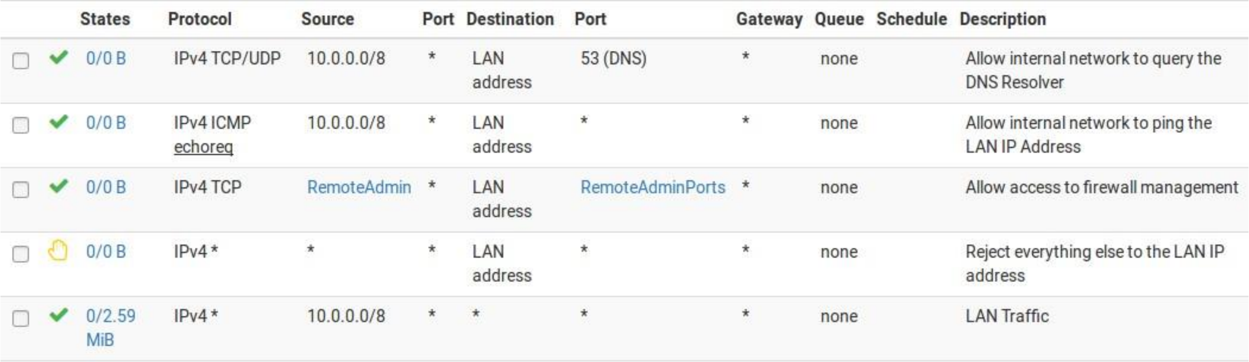

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
□	✓	0/0 B	IPv4 TCP/UDP	10.0.0.0/8	*	LAN address	53 (DNS)	*	none		Allow internal network to query the DNS Resolver
□	✓	0/0 B	IPv4 ICMP echoreq	10.0.0.0/8	*	LAN address	*	*	none		Allow internal network to ping the LAN IP Address
□	✓	0/0 B	IPv4 TCP	RemoteAdmin	*	LAN address	RemoteAdminPorts	*	none		Allow access to firewall management
□	✋	0/0 B	IPv4 *	*	*	LAN address	*	*	none		Reject everything else to the LAN IP address
□	✓	0/2.59 MiB	IPv4 *	10.0.0.0/8	*	*	*	*	none		LAN Traffic

Fig. 10: Example Restricted Management LAN Rules

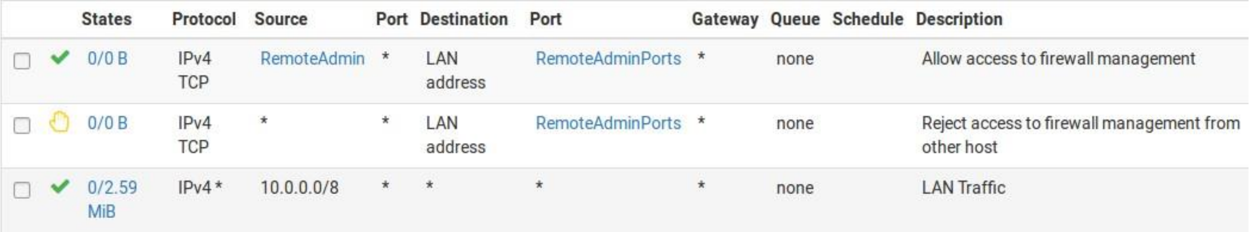

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
□	✓	0/0 B	IPv4 TCP	RemoteAdmin	*	LAN address	RemoteAdminPorts	*	none		Allow access to firewall management
□	✋	0/0 B	IPv4 TCP	*	*	LAN address	RemoteAdminPorts	*	none		Reject access to firewall management from other host
□	✓	0/2.59 MiB	IPv4 *	10.0.0.0/8	*	*	*	*	none		LAN Traffic

Fig. 11: Restricted Management LAN Rules Alternate Example

Once the firewall rules are configured, disable the webGUI anti-lockout rule on the System > Advanced page (Figure *Anti-Lockout Rule Disabled*). Check the box and click Save.

Note: If the management interface can no longer be accessed after disabling the anti-lockout rule, the firewall rules were not configured appropriately. Re-enable the anti-lockout rule by using the Set Interface(s) IP address option at the console menu, then choose to reset the LAN IP address. Set it to its current IP address, and the rule will automatically be re-enabled.

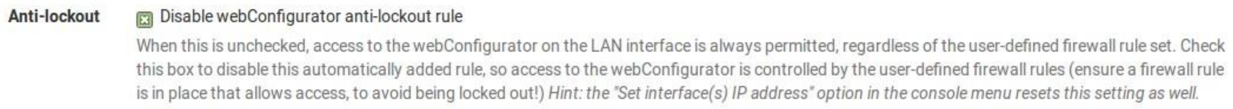

Anti-lockout ☒ Disable webConfigurator anti-lockout rule

When this is unchecked, access to the webConfigurator on the LAN interface is always permitted, regardless of the user-defined firewall rule set. Check this box to disable this automatically added rule, so access to the webConfigurator is controlled by the user-defined firewall rules (ensure a firewall rule is in place that allows access, to avoid being locked out!) *Hint: the "Set interface(s) IP address" option in the console menu resets this setting as well.*

Fig. 12: Anti-Lockout Rule Disabled

Anti-spoofing Rules

AZTCO-FW uses the antispoof feature in pf to block spoofed traffic. This provides Unicast Reverse Path Forwarding (uRPF) functionality as defined in RFC 3704. The firewall checks each packet against its routing table, and if a connection attempt comes from a source IP address on an interface where the firewall knows that network does not reside, it is dropped. For example, a packet coming in WAN with a source IP address of an internal network is dropped. Anything initiated on the internal network with a source IP address that does not reside on the internal network is dropped.

Block Private Networks

The Block private networks option on the WAN interface automatically puts in a block rule for RFC 1918 subnets. Unless private IP space is in use on the WAN, enable this option. This only applies to traffic initiated on the WAN side. Local clients may still reach hosts on private networks from the inside of the firewall. This option is available for any interface, but is generally only used on WAN type interfaces. A similar rule can be created manually to block private networks on interfaces by creating an alias containing the RFC 1918 subnets and adding a firewall rule to the top of the interface rules to block traffic with a source matching that alias. (See *Private IP Addresses* for more information about private IP addresses.)

Block Bogon Networks

Bogon networks are those which should never be seen on the Internet, including reserved and unassigned IP address space. The presence of traffic from these networks can indicate either spoofed traffic or an unused subnet that has been hijacked for malicious use. Bogon lists are intended to filter invalid traffic from the Internet (e.g. on WANs) coming to the firewall for cases where the source cannot be otherwise filtered or validated, such as for public services. If rules on an interface only allow from specific remote sources, bogon blocking does not offer any benefit. AZTCO-FW software provides two bogons lists that are updated as needed, one for IPv4 bogon networks and one for IPv6 bogon networks.

Warning: Blocking bogon networks is not suited for use on local/private interfaces such as LAN. Blocking bogon networks on local interfaces can be harmful as they will block traffic which is necessary for proper local network operations, especially for IPv6. If local interfaces have proper rules which only allow from specific local sources, bogon blocking is unnecessary.

The firewall fetches an updated bogons list on the first day of each month from AZTCO-FW servers. The script runs at 3:00 a.m. local time, and sleeps a random amount of time up to 12 hours before performing the update. This list does not change frequently, and new IP address assignments are removed from the bogons list months before they are allocated, so a monthly update is adequate. To automatically update the list more frequently, change the Update Frequency for bogons under System > Advanced on the Firewall & NAT tab.

Note: The bogons list for IPv6 is quite large, and may not load if there is not enough memory in the system, or if the maximum number of table entries is not large enough to contain it. See *Firewall Maximum Table Entries* for information on changing that value.

See also:

For information on troubleshooting bogon updates and forcing manual updates, see *Troubleshooting Bogon Network List Updates*.

IPsec

When a site to site IPsec connection is enabled, rules are automatically added allowing the remote tunnel endpoint IP address access to UDP ports 500 and 4500, and the ESP protocol on the WAN IP address used for the connection. When IPsec for mobile clients is enabled the same traffic is allowed, but from a source of *any*, rather than a specific source address.

Because of the way policy routing works, any traffic that matches a rule specifying a gateway will be forced out to the Internet and will bypass IPsec processing. Rules are added automatically to negate policy routing for traffic destined to remote VPN subnets, but they do not always have the intended effect. To disable the automatic negation rules, see *Disable Negate rules* and add a firewall rule at the *top* of the rules on the internal interface to pass traffic to the VPN *without a gateway set*.

See also:

Automatically added IPsec rules are discussed in further depth in *IPsec*.

Default Deny Rule

Rules that do not match any user-defined rules nor any of the other automatically added rules are silently blocked by the default deny rule (as discussed in *Default Deny*).

Configuring firewall rules

When configuring firewall rules in the AZTCO-FW® WebGUI under Firewall > Rules many options are available to control how traffic is matched and controlled. Each of these options are listed in this section.

Action

This option specifies whether the rule will *pass*, *block*, or *reject* traffic.

Pass A packet matching this rule will be allowed to pass through the firewall. If state tracking is enabled for the rule, a state table entry is created which allows related return traffic to pass back through. See *Stateful Filtering* for more information.

Block A packet matching this rule will be discarded.

Reject A packet matching this rule will be discarded and for supported protocols, a message will be sent back to the originator indicating that the connection was refused.

See also:

See *Block vs. Reject* for a deeper description of the options and for help deciding between Block and Reject.

Disabled

To disable a rule without removing it from the rule list, check this box. It will still show in the firewall rules screen, but the rule will appear grayed out to indicate its disabled state.

Interface

The Interface drop down specifies the interface receiving traffic to be controlled by this rule. Remember that on interface and group tab rules, traffic is only filtered on the interface where the traffic is *initiated*. Traffic initiated from the LAN destined to the Internet or any other interface on the firewall is filtered by the *LAN* ruleset.

TCP/IP Version

Instructs the rule to apply for *IPv4*, *IPv6*, or both *IPv4+IPv6* traffic. The rules will only match and act upon packets matching the correct protocol. Aliases may be used which contain both types of IP addresses and the rule will match only the addresses from the correct protocol.

Protocol

The protocol this rule will match. Most of these options are self-explanatory. *TCP/UDP* will match both *TCP* and *UDP* traffic. Specifying *ICMP* will show an additional drop down box to select the ICMP type. Several other common protocols are also available.

Note: This field defaults to *TCP* for a new rule because it is a common default and it will display the expected fields for that protocol. To make the rule apply to any protocol, change this field to *any*. One of the most common mistakes in creating new rules is accidentally creating a TCP rule and then not being able to pass other non-TCP traffic such as ping, DNS, etc.

ICMP Type

When *ICMP* is selected as the protocol, this drop-down contains all possible ICMP types to match. When passing ICMP, the best practice is to only pass the required types when feasible. The most common use case is to pass only a type of *Echo Request* which will allow an ICMP ping to pass.

Tip: Historically, ICMP has a bad reputation but it is generally beneficial and does not deserve the reputation on modern networks. Allowing an ICMP type of *any* is typically acceptable when allowing ICMP.

Source

This field specifies the source IP address, subnet, or alias that will match this rule.

The drop-down box for source allows several different pre-defined types of sources:

Any Matches any address.

Single host or Alias Matches a single IP address or alias name. When this is active, an alias name may be typed in the Source Address field.

Network Uses both an IP address and subnet mask to match a range of addresses.

PPPoE Clients A macro that will match traffic from the client address range for the PPPoE server if the PPPoE server is enabled.

L2TP Clients A macro that will match traffic from the client address range for the L2TP server if the L2TP server is enabled.

Interface Net An entry in this list is present for each interface on the firewall. These macros specify the subnet for that interface exactly, including any IP alias VIP subnets that differ from the defined interface subnet.

Interface Address An entry in this list is present for each interface on the firewall. These macros specify the IP address configured on that interface.

Warning: The *WAN Net* choice for source or destination means the subnet of the WAN interface only. It does not mean “The Internet” or any remote host.

For rules matching TCP and/or UDP, the source port may also be specified by clicking the ⚙ Display Advanced. The source port is hidden behind the Display Advanced button because normally the source port must remain set to *any*, as TCP and UDP connections are sourced from a random port in the ephemeral port range (between 1024 through 65535, the exact range used varying depending on the OS and OS version that is initiating the connection). The source port is almost never the same as the destination port, and it should never be configured as such unless the application in use is known to employ this atypical behavior. It is also safe to define a source port as a range from 1024 to 65535.

Selecting Invert Match will negate the match so that all traffic except this source value will trigger the rule.

Destination

This field specifies the destination IP address, subnet, or alias that will match this rule. See the description of the Source option in *Source* for more details.

For rules specifying TCP and/or UDP, the destination port, port range, or alias is also specified here. Unlike source, configuring a destination port is required in many cases, as it is more secure than using *any* and usually the destination port will be known in advance based on the protocol. Many common port values are available in the drop-down lists, or select *(other)* to enter a value manually or to use a port alias.

Tip: To specify a continuous range of ports, enter the lower port in the From section and the higher port value in the To section.

Log

This box determines whether packets that match this rule will be logged to the firewall log. Logging is discussed in more detail in *Logging Practices*.

Description

Enter a description here for reference. This is optional, and does not affect functionality of the rule. The best practice is to enter text describing the purpose of the rule. The maximum length is 52 characters.

Advanced Options

Options which are less likely to be required or that have functionality confusing to new users have been tucked away in this section of the page. Click ⚙ Display Advanced to show all of the advanced options. If an option in this section of the page has been set, then it will appear when the rule is loaded in the future .

Source OS

One of the more unique features of pf and thus AZTCO-FW is the ability to filter by the operating system initiating a connection. For TCP rules, pf enables passive operating system fingerprinting (“p0f”) that allows rules to match based on the operating system initiating the TCP connection. The p0f feature of pf determines the OS in use by comparing characteristics of the TCP SYN packet that initiates TCP connections with a fingerprints file. Note that it is possible to change the fingerprint of an operating system to look like another OS, especially with open source operating systems such as the BSDs and Linux. This isn’t easy, but if a network contains technically proficient users with administrator or root level access to systems, it is possible.

Diffserv Code Point

Differentiated Services Code Point is a way for applications to indicate inside the packets how they would prefer routers to treat their traffic as it gets forwarded along its path. The most common use of this is for quality of service or traffic shaping purposes. The lengthy name is often shortened to *Diffserv Code Point* or abbreviated as *DSCP* and sometimes referred to as the *TOS field*.

The program or device generating the packets, for example Asterisk via its tos_sip and tos_audio configuration parameters, will set the DSCP field in the packets and then it is up to the firewall and other interim routers to match and queue or act on the packets.

To match these parameters in the firewall, use the Diffserv Code Point drop-down entry that matches the value set by the originating device. There are numerous options, each with special meaning specific to the type of traffic. Consult the documentation for the device originating the traffic for more detail on which values must be matched.

The downside of DSCP is that it assumes routers support or act on the field, which may or may not be the case. Different routers may treat the same DSCP value in unintended or mismatched ways. Worse yet, some routers will clear the DSCP field in packets entirely as it forwards them. Also, the way pf matches traffic, the DSCP value must be set on the first packet of a connection creating a state, as each packet is not inspected individually once a state has been created.

Note: This option only reads and matches the DSCP value. It does not set a value in packets.

IP Options

Checking this box will allow packets with defined IP options to pass. By default, pf blocks all packets that have IP options set in order to deter OS fingerprinting, among other reasons. Check this box to pass IGMP or other multicast traffic containing IP options.

Disable Reply-To

The firewall adds the reply-to keyword to rules on WAN type interfaces by default to ensure that traffic that enters a WAN will also leave via that same WAN. In certain cases this behavior is undesirable, such as when some traffic is routed via a separate firewall/router on the WAN interface. In these cases, check this option to disable reply-to only for traffic matching this rule, rather than disabling reply-to globally.

Tag and Tagged

The Tag and Tagged fields are useful in concert with floating rules, so the firewall can mark a packet with a specific string as it enters an interface, and then act differently on a matched packet on the way out with a floating rule. See *Marking and Matching* for more on this topic.

Maximum state entries this rule can create

This option limits the maximum number of connections, total, that can be allowed by this rule. If more connections match this rule while it is at its connection limit, this rule will be skipped in the rule evaluation. If a later rule matches, the traffic has the action of that rule applied, otherwise it hits the default deny rule. Once the number of connections permitted by this rule drops below this connection limit, traffic can once again match this rule.

Maximum number of unique source hosts

This option specifies how many total source IP addresses may simultaneously connect for this rule. Each source IP address is allowed an unlimited number of connections, but the total number of distinct source IP addresses allowed is restricted to this value.

Maximum number of established connections per host

To limit access based on connections per host, use this setting. This value can limit a rule to a specific number of connections per source host (e.g. 10), instead of a specific global connection total. This option controls how many fully established (completed handshake) connections are allowed per host that match the rule. This option is only available for use with TCP connections.

Maximum state entries per host

This setting works similar to the established count above, but it checks for state entries alone rather than tracking if a successful connection was made.

Maximum new connections / per second

This method of rate limiting helps ensure that a high TCP connection rate will not overload a server or the state table on the firewall. For example, limits can be placed on incoming connections to a mail server, reducing the burden of being overloaded by spambots. It can also be used on outbound traffic rules to set limits that would prevent any single machine from loading up the state table on the firewall or making too many rapid connections, behaviors which are common with viruses. A connection amount and a number of seconds for the time period may be configured for the rule. Any IP address exceeding the specified number of connections within the given time frame will be blocked by the firewall for one hour. Behind the scenes, this is handled by the virusprot table, named for its typical purpose of virus protection. This option is only available for use with TCP connections.

State timeout in seconds

Using this field, a state timeout for traffic matching this rule may be defined, overriding the default state timeout. Any inactive connections will be closed when the connection has been idle for this amount of time. The default state timeout depends on the firewall optimization algorithm in use. The optimization choices are covered in *Firewall Optimization Options*

Note: This option only controls the traffic in the inbound direction, so it is not very useful on its own. Outbound traffic for a matching connection will still have the default state timeout. To use this setting properly, a matching floating rule is also required in the outbound path taken by the traffic with a similar state timeout setting.

TCP Flags

By default, new pass rules for TCP only check for the TCP SYN flag to be set, out of a possible set of SYN and ACK. To account for more complex scenarios, such as working around asymmetric routing or other non-traditional combinations of traffic flow, use this set of controls to change how the flags are matched by the firewall rule.

The first row controls which flags must be set to match the rule. The second row defines the list of flags that will be consulted on the packet to look for a match.

The meanings of the most commonly used flags are:

SYN Synchronize sequence numbers. Indicates a new connection attempt.

ACK Indicates ACKnowledgment of data. These are replies to let the sender know data was received OK.

FIN Indicates there is no more data from the sender, closing a connection.

RST Connection reset. This flag is set when replying to a request to open a connection on a port which has no listening daemon. Can also be set by firewall software to turn away undesirable connections.

PSH Indicates that data should be pushed or flushed, including data in this packet, by passing the data up to the application.

URG Indicates that the urgent field is significant, and this packet should be sent before data that is not urgent.

To allow TCP with any flags set, check Any Flags.

State Type

There are three options for state tracking in AZTCO-FW that can be specified on a per-rule basis:

Keep When chosen, the firewall will create and maintain a state table entry for permitted traffic. This is the default, and the best choice in most situations.

Sloppy State Sloppy is a less strict means of keeping state that is intended for scenarios with asymmetric routing. When the firewall can only see half the traffic of a connection, the validity checks of the default state keeping will fail and traffic will be blocked. Mechanisms in pf that prevent certain kinds of attacks will not kick in during a sloppy state check.

Synproxy This option causes AZTCO-FW to proxy incoming TCP connections. TCP connections start with a three way handshake. The first packet of a TCP connection is a SYN from source, which elicits a SYN ACK response from the destination, then an ACK in return from the source to complete the handshake. Normally the host behind the firewall will handle this on its own, but synproxy state has

the firewall complete this handshake instead. This helps protect against one type of Denial of Service attack, SYN floods. This is typically only used with rules on WAN interfaces. This type of attack is best handled at the target OS level today, as every modern operating system includes capabilities of handling this on its own. Because the firewall can't know what TCP extensions the back-end host supports, when using synproxy state, it announces no supported TCP extensions. This means connections created using synproxy state will not use window scaling, SACK, nor timestamps which will lead to significantly reduced performance in most all cases. It can be useful when opening TCP ports to hosts that do not handle network abuse well, where top performance isn't a concern.

None This option will not keep state on this rule. This is only necessary in some highly specialized advanced scenarios, none of which are covered in this documentation because they are exceedingly rare.

Note: Setting *None* here only affects traffic in the inbound direction, so it is not very useful on its own since a state will still be created in the outbound direction. It must be paired with a floating rule in the outbound direction which also has the same option chosen.

No XML-RPC Sync

Checking this box prevents this rule from synchronizing to other High Availability cluster members via XMLRPC. This is covered in *High Availability*. This does not prevent a rule on a secondary node from being overwritten by the primary.

VLAN Priority (Match and Set)

802.1p, also known as IEEE P802.1p or Priority Code Point, is a way to match and tag packets with a specific quality of service priority. Unlike DSCP, 802.1p operates at layer 2 with VLANs. However, like DSCP, the upstream router must also support 802.1p for it to be useful.

There are two options in this section. The first will match an 802.1p field so the firewall can act on it. The second will inject an 802.1p tag into a packet as it passes through this firewall. Some ISPs may require an 802.1p tag to be set in certain areas, such as France, in order to properly handle voice/video/data on segregated VLANs at the correct priority to ensure quality.

There are eight levels of priority for 802.1p, and each has a two letter code in the GUI. In order from lowest priority to highest, they are:

BK Background

BE Best Effort

EE Excellent Effort

CA Critical Applications

VI Video

VO Voice

IC Internetwork Control

NC Network Control

Schedule

This option configures a schedule specifying the days and times for the rule to be in effect. Selecting “none” means the rule will always be enabled. For more information, see *Time Based Rules* later in this chapter.

Gateway

This option configures a Gateway or Gateway Group to be used by traffic matching this rule. This is covered in *Policy routing*.

In/Out Pipe (Limiters)

These selections list defined Limiters to apply a bandwidth limit to the traffic entering this interface (In) and leaving this interface (Out). More detail on limiters can be found in *Limiters*.

Ackqueue/Queue

These options define which ALTQ traffic shaper queues are applied to traffic entering and exiting this interface. For more information on traffic shaping, see *Traffic Shaper*.

Floating Rules

Floating Rules are a special type of advanced rule that can perform complicated actions not possible with rules on interface or group tabs. Floating rules can act on multiple interfaces in the inbound, outbound, or both directions. The use of inbound and outbound filtering makes designing the rules more complex and prone to user error, but they can be desirable in specific applications.

Most firewall configurations will never have floating rules, or only have them from the traffic shaper.

Precautions/Caveats

Floating rules can be a lot more powerful than other rules, but also more confusing, and it is easier to make an error that could have unintended consequences in passing or blocking traffic.

Floating rules in the inbound direction, applied to multiple WANs, will not get reply-to added as they would with individual interface rules, so the same problem exists here as existed with interface groups: The traffic will always exit the WAN with the default gateway, and not return properly out the WAN it entered.

Given the relative unfamiliarity of many users with Floating rules, they may not think to look there for rules when maintaining the firewall. As such, they can be a little more difficult for administration since it may not be an obvious place to look for rules.

Be careful when considering the source and destination of packets depending on the inbound and outbound direction. For example, rules in the outbound direction on a WAN would have a local source of the firewall (after NAT) and remote destination.

Potential Uses

The most common use of Floating rules is for ALTQ traffic shaping. Floating tab rules are the only type of rules which can match and queue traffic without explicitly passing the traffic.

Another way to use floating rules is to control traffic leaving from the firewall itself. Floating rules can prevent the firewall from reaching specific IP addresses, ports, and so on.

Other common uses are to ensure that no traffic can exit from other paths into a secure network, no matter what rules exist on other interfaces. By blocking outbound toward a secure network from all but the approved locations, the likelihood of later accidentally allowing traffic in through some other unintended path is reduced. Similarly, they can be used to prevent traffic destined for private networks from leaving a WAN interface, to prevent VPN traffic from leaking.

As mentioned earlier in the interface rules, they can also effectively enact state timeouts, tag/match operations, “no state” rules, and “sloppy state” rules for asymmetric routing.

Processing Order

In the inbound direction, floating rules work essentially the same as interface or group rules except that they are processed first. In the outbound direction, however, things get a little more confusing.

Firewall rules are processed after NAT rules, so rules in the outbound direction on a WAN can never match a local/private IP address source if outbound NAT is active on that interface. By the time it hits the rule, the source address of the packet is now the WAN interface IP address. In most cases this can be worked around by using the match options to tag a packet on the LAN on the way in and then matching that tag on the way out of the firewall.

Floating rules are processed before interface group rules and interface rules, so that must also be taken into consideration.

Match Action

The *match* action is unique to floating rules. A rule with the *match* action will not pass or block a packet, but only match it for purposes of assigning traffic to queues or limiters for traffic shaping. Match rules do not work with *Quick* enabled.

Quick

Quick controls whether rule processing stops when a rule is matched. The Quick behavior is added to all interface tab rules automatically, but on floating rules it is optional. Without *Quick* checked, the rule will only take effect if no other rules match the traffic. It reverses the behavior of “first match wins” to be “last match wins”.

Using this mechanism, a default action of sorts can be crafted which will take effect only when no other rules match, similar to the default block rules on WANs.

In most situations, we advise having Quick selected. There are certain specific scenarios where leaving Quick unchecked is necessary, but they are few and far between. For most scenarios, the only rules they would have without quick selected are *match* rules traffic shaper rules.

Interface

The Interface selection for floating rules is different than the one for normal interface rules: It is a multi-select box so one, multiple, or all possible interfaces may be selected. Ctrl-click on interfaces to select them one by one, or use other combinations of click/drag or shift-click to select multiple interfaces.

Direction

Floating rules are not limited to the inbound direction like interface rules. They can also act in the outbound direction by selecting *out* here, or in both directions by selecting *any*. The *in* direction is also available.

The *out* direction is useful for filtering traffic from the firewall itself, for matching other undesirable traffic trying to exit an interface, or for fully configuring “sloppy state” rules, “no state” rules, or alternate state timeouts.

Marking and Matching

Using the Tag and Tagged fields, a connection can be marked by an interface tab rule and then matched in the outbound direction on a floating rule. This is a useful way to act on WAN outbound traffic from one specific internal host that could not otherwise be matched due to NAT masking the source. It can also be used similarly for applying shaping outbound on WAN from traffic specifically tagged on the way into the firewall.

For example, on a LAN rule, use a short string in the Tag field to mark a packet from a source of 10.3.0.56. Then on a floating rule, quick, outbound on WAN, use Tagged with the same string to act on the traffic matched by the LAN rule.

Time Based Rules

Time based rules allow firewall rules to activate during specified days and/or time ranges. Time based rules function the same as any other rule, except they are effectively not present in the ruleset outside of their scheduled times.

Time Based Rules Logic

When dealing with time-based rules, the schedule determines when to apply the action specified in the firewall rule. When the current time or date is not covered by the schedule, the firewall acts as if the rule is not there. For example, a rule that passes traffic on Saturdays will only block it on other days if a separate block rule exists underneath it. The rules are processed from the top-down, the same as other firewall rules. The first match is used, and once a match is found, that action is taken if the rule is in schedule, and no other rules are evaluated.

Tip: Remember when using schedules that the rule will have no effect outside of their scheduled times. The rule will not have its action reversed because the current time is not within the scheduled time. Failing to account for this behavior could result in giving clients unintended access outside of the defined time ranges in a schedule.

Configuring Schedules for Time Based Rules

Schedules must be defined before they can be used on firewall rules. Schedules are defined under Firewall > Schedules, and each schedule can contain multiple time ranges. In the following example, a company wants to deny access to HTTP during business hours, and allow it all other times of the day.

Defining Times for a Schedule

To add a schedule:

- Navigate to Firewall > Schedules
- Click ➕ Add to bring up the schedule editing screen, as seen in Figure *Adding a Time Range*.
- Enter a Schedule Name. This is the name that will appear in the selection list for use in firewall rules. Much like alias names, this name must only contain letters and digits, no spaces. For example: BusinessHours

- Enter a Description of this schedule, such as Normal Business Hours.
- Define one or more time ranges:
  - Set the Month by selecting a specific month and days, or by clicking the day of the week header for weekly recurring schedules.
  - Choose a Start Time and Stop Time which control when the rule is active on the selected days. The time cannot cross midnight on any day. A full day is *0:00* to *23:59*.
  - Enter an optional Time Range Description for this specific range, e.g. Work Week
  - Click + Add Time to add the choice as a range
  - Repeat Month, Time, and + steps for additional ranges
- Click Save

A schedule can apply to specific days, such as September 2, 2016, or to days of the week, such as Monday-Wednesday. To select any given day within the next year, choose the Month from the drop-down list, then click on the specific day or day numbers on the calendar. To select a day of the week, click its name in the column headers.

For this example, click on Mon, Tue, Wed, Thu, and Fri. This will make the schedule active for any Monday-Friday, regardless of the month. Now select the time for this schedule to be active, in 24-hour format. The hours for this example business are *9:00* to *17:00* (5pm). All times are given in the local time zone.

Once the time range has been defined, it will appear in the list at the bottom of the schedule editing screen, as in Figure *Added Time Range*.

To expand on this setup, there may be a half day on Saturday to define, or maybe the shop opens late on Mondays. In that case, define a time range for the identical days, and then another range for each day with different time ranges. This collection of time ranges will be the full schedule.

Once the schedule entry has been saved, the browser will return to the schedule list, as in Figure *Schedule List After Adding*. This schedule will now be available for use in firewall rules.

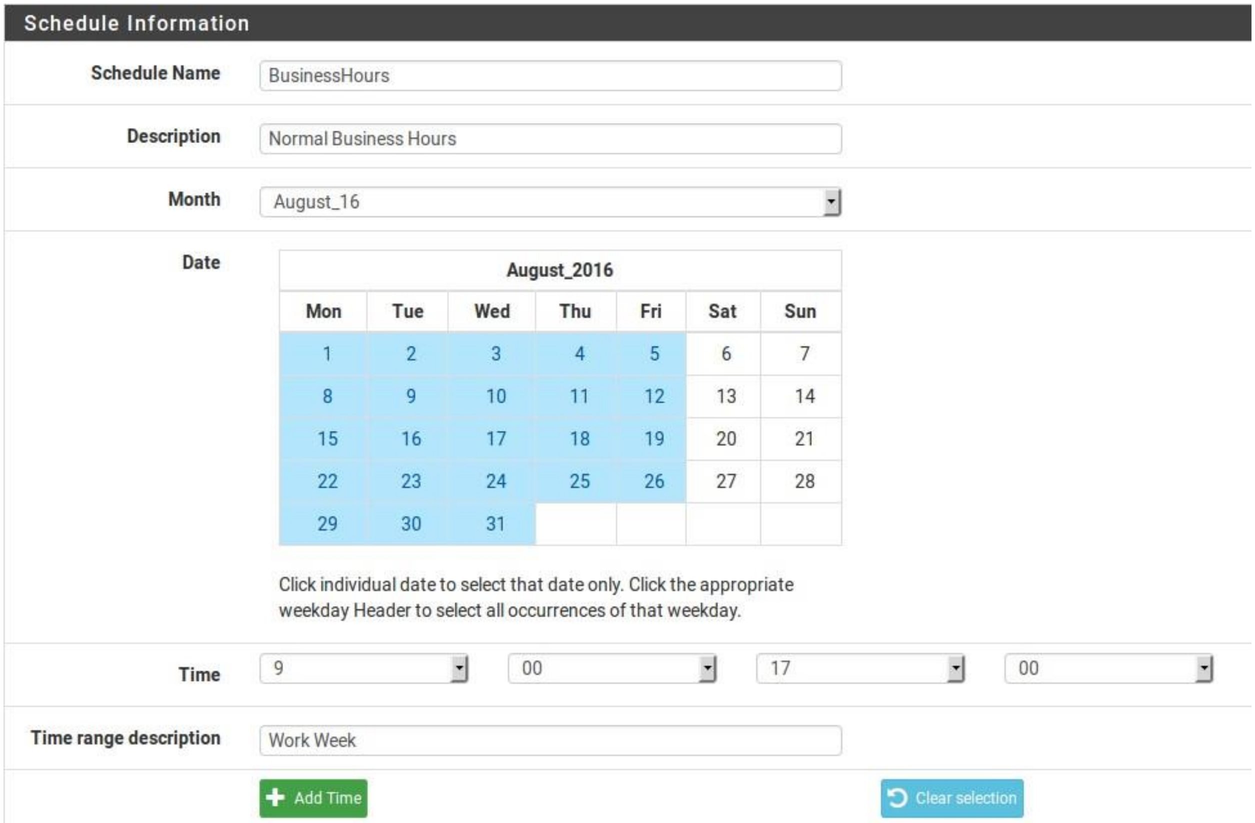

Schedule Information

Schedule Name	BusinessHours
Description	Normal Business Hours
Month	August_16

Date

August_2016

Mon	Tue	Wed	Thu	Fri	Sat	Sun
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	31				

Click individual date to select that date only. Click the appropriate weekday Header to select all occurrences of that weekday.

Time	9	00	17	00
Time range description	Work Week			

Fig. 13: Adding a Time Range

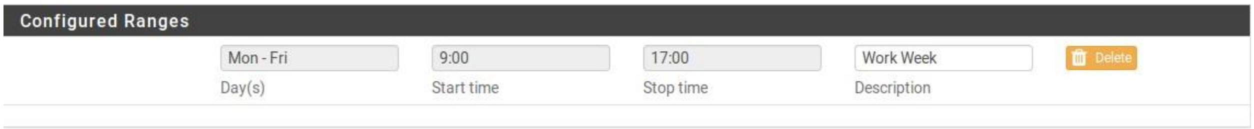

Configured Ranges

Day(s)	Start time	Stop time	Description	
Mon - Fri	9:00	17:00	Work Week	Delete

Fig. 14: Added Time Range

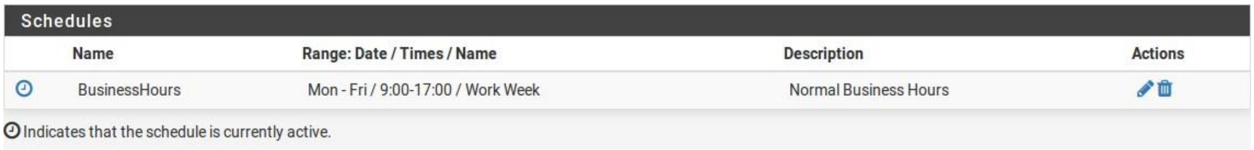

Schedules

	Name	Range: Date / Times / Name	Description	Actions
⏲	BusinessHours	Mon - Fri / 9:00-17:00 / Work Week	Normal Business Hours	✎🗑

⏲ Indicates that the schedule is currently active.

Fig. 15: Schedule List After Adding **Using the Schedule in a Firewall Rule**

To create a firewall rule employing this schedule, create a new rule on the desired interface. See *Adding a firewall rule* and *Configuring firewall rules* for more information about adding and editing rules. For this example, add a rule to reject TCP traffic on the LAN interface from the LAN subnet to any destination on the HTTP port. In the advanced options for the rule, locate the Schedule setting and choose the *BusinessHours* schedule, as in Figure *Choosing a Schedule for a Firewall Rule*.

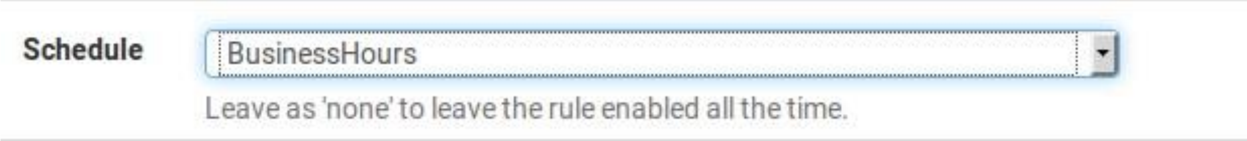

Schedule BusinessHours

Leave as 'none' to leave the rule enabled all the time.

Fig. 16: Choosing a Schedule for a Firewall Rule

After saving the rule, the schedule will appear in the firewall rule list along with an indication of the schedule's active state. As shown in Figure *Firewall Rule List with Schedule*, this is a reject rule, and the schedule column indicates that the rule is currently in its active blocking state because it is being viewed at a time within the scheduled range. If the mouse cursor hovers over the schedule state indicator, a tooltip is displayed by the firewall showing how the rule will behave at the current time. Since this is being viewed inside of the times defined in the BusinessHours schedule, this will say "Traffic matching this rule is currently being denied". If there is a pass rule that would match the traffic out on port 80 from the LAN net after this rule, then it would be allowed outside of the scheduled hours.

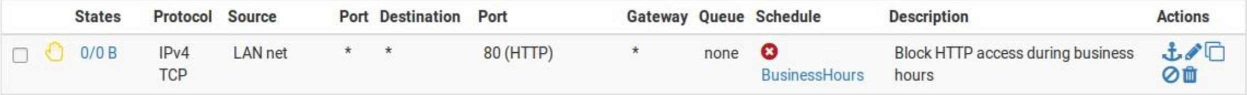

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
☐	✋	0/0 B	IPv4 TCP	LAN net	*	*	80 (HTTP)	*	none	✖ BusinessHours	Block HTTP access during business hours	⚓ ✎ ❐ ⊘ 🗑

Fig. 17: Firewall Rule List with Schedule

Now that the rule is defined, test it both inside and outside of the scheduled times to ensure that the desired behavior is enacted.

Tip: By default, states are cleared for active connections permitted by a scheduled rule when the schedule expires. This shuts down access for anyone allowed by the rule while it was active. To allow these connections to remain open, check Do not kill connections when schedule expires under System > Advanced on the Miscellaneous tab.

Viewing the pf ruleset

AZTCO-FW® software handles translating the firewall rules in the GUI into a set of rules which can be interpreted by the packet filter (PF).

Generated Rules

The PF rules generated by the firewall configuration are in /tmp/rules.debug. However, that file cannot be edited to make persistent changes - it will be overwritten by the next filter reload event.

Note: There is rarely a need to manually edit firewall rules generated by the GUI. In most cases if it appears to be necessary, something is incorrect with the configuration.

If the generated rules truly must be edited, then the edits must be made to the source code which generates the ruleset in /etc/inc/filter.inc. Such changes will be lost when updating to a new version.

Interpreted Rules

PF can interpret the rules slightly differently than in the way they were generated by the filter code. To view the rule set as has been interpreted by PF, use one of the following methods.

Using the *SSH console* or *Command Prompt* field in the GUI, run the following:

Show Firewall Rules:

```
# pfctl -sr
```

Show NAT rules:

```
# pfctl -sn
```

Show all:

```
# pfctl -sa
```

For more verbose output including rule counters, ID numbers, and so on, use:

```
# pfctl -vvsr
```

There may be additional rules in anchors from packages or features such as UPnP. To view these rules, use:

```
# pfSsh.php playback pfanchordrill
```

Methods of Using Additional Public IP Addresses

Methods of deploying additional public IP addresses vary depending on how the addresses are delegated, the size of the allocation, and the goals for the specific network environment. To use additional public IP addresses with NAT, for example, the firewall will need *Virtual IP Addresses*.

There are two options for directly assigning public IP addresses to hosts: Routed public IP subnets and bridging.

Choosing between routing, bridging, and NAT

Additional public IP addresses can be put to use by directly assigning them on the systems that will use them, or by using NAT. The available options depend on how the addresses are allocated by the ISP.

Additional static IP addresses

Methods of using additional static public IP addresses vary depending on the type of assignment. Each of the common scenarios is described here.

Single IP Subnet on WAN

With a single public IP subnet on WAN, one of the public IP addresses will be on the upstream router, commonly belonging to the ISP, and another one of the IP addresses will be assigned as the WAN IP address on AZTCO-FW® software. The remaining IP addresses can be used with either NAT, bridging or a combination of the two.

To use the addresses with NAT, add Proxy ARP, IP alias or CARP type Virtual IP addresses.

To assign public IP addresses directly to hosts behind the firewall, a dedicated interface for those hosts must be bridged to WAN. When used with bridging, the hosts with the public IP addresses directly assigned must use the same default gateway as the WAN of the firewall: the upstream ISP router. This will create difficulties if the hosts with public IP addresses need to initiate connections to hosts behind other interfaces of the firewall, since the ISP gateway will not route traffic for internal subnets back to the firewall.

Figure *Multiple Public IP addresses In Use Single IP Subnet* shows an example of using multiple public IP addresses in a single block with a combination of NAT and bridging.

See also:

For information on configuration, NAT is discussed further in *Network Address Translation*, and bridging in *Bridging*.

Small WAN IP Subnet with Larger LAN IP Subnet

Some ISPs will allocate a small IP subnet as the "WAN side" assignment, sometimes called a transport or interconnect network, and route a larger "inside" subnet to the firewall. Commonly this is a /30 on the WAN side and a /29 or larger for use inside the firewall. The service provider router is assigned one end of the /30, typically the lowest IP address, and the firewall is assigned the higher IP address. The provider then routes the second subnet to the WAN IP address of the firewall. The additional IP subnet may be used by the firewall on a routed LAN or OPT interface with public IP addresses directly assigned to hosts, with NAT using Other type VIPs, or a combination of the two. Since the IP addresses are routed to the firewall, ARP is not needed so VIP entries are not necessary for use with NAT.

Because AZTCO-FW is the gateway on the local segment, routing from the public local subnet hosts to LAN is much easier than in the bridged scenario required when using a single public IP subnet. Figure *Multiple Public IP Addresses Using Two IP Subnets* shows an example that combines a routed IP subnet and NAT. Routing public IP addresses is covered in *Routing Public IP Addresses*, and NAT in *Network Address Translation*.

If the firewall is part of a High Availability cluster using CARP, the WAN side subnet will need to be a /29 so each firewall has its own WAN IP address plus a CARP VIP. The provider will route the larger inside subnet to the WAN CARP VIP in this type of configuration. The inside IP subnet must be routed to an IP address that is always available regardless of which firewall is up, and the smallest subnet usable with CARP is a /29. Such a setup with CARP is the same as illustrated above, with the OPT1 gateway being a CARP VIP, and the provider routing to a CARP VIP rather than the WAN IP address. CARP is covered in *High Availability*.

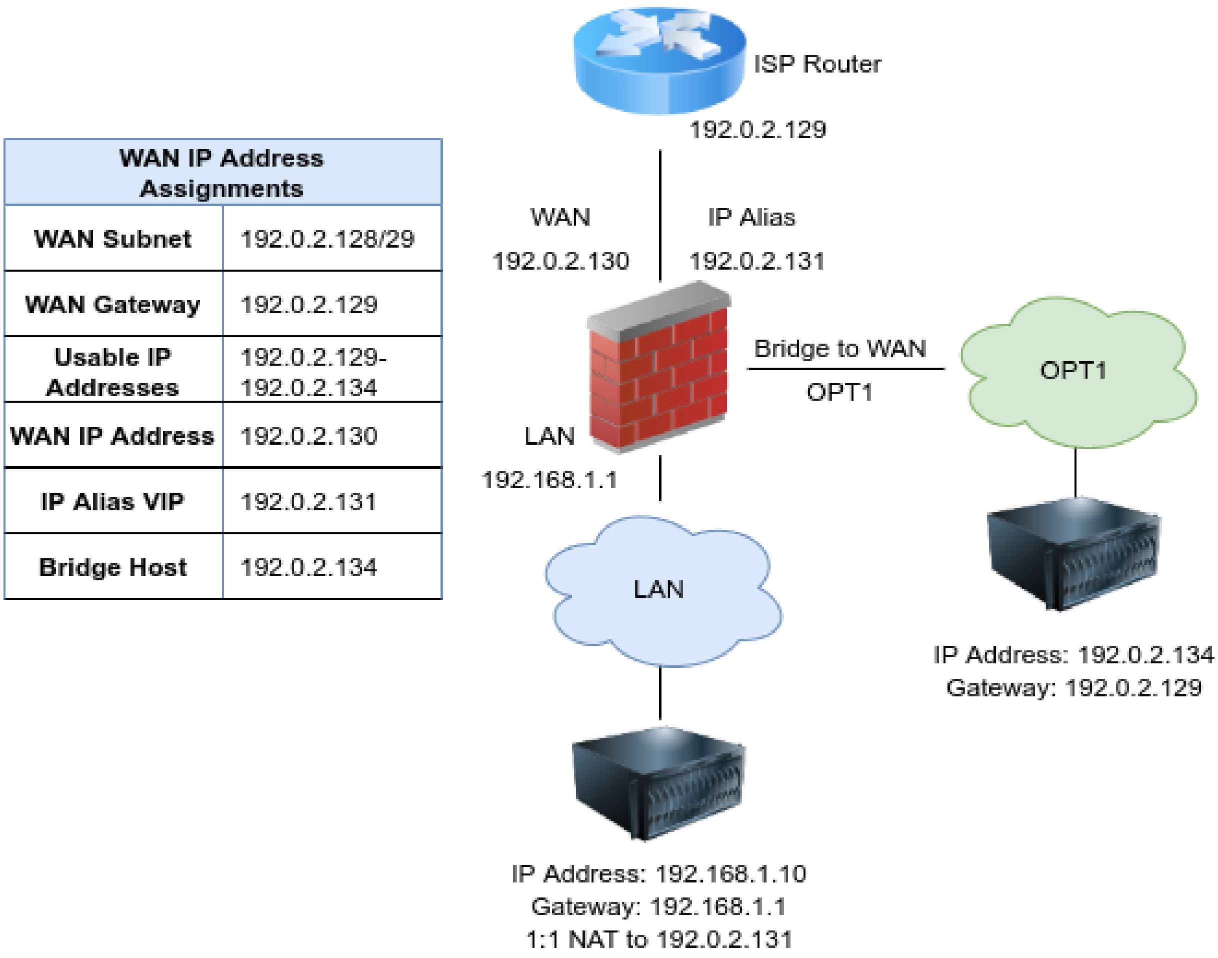

WAN IP Address Assignments	
WAN Subnet	192.0.2.128/29
WAN Gateway	192.0.2.129
Usable IP Addresses	192.0.2.129- 192.0.2.134
WAN IP Address	192.0.2.130
IP Alias VIP	192.0.2.131
Bridge Host	192.0.2.134

Fig. 18: Multiple Public IP addresses In Use Single IP Subnet

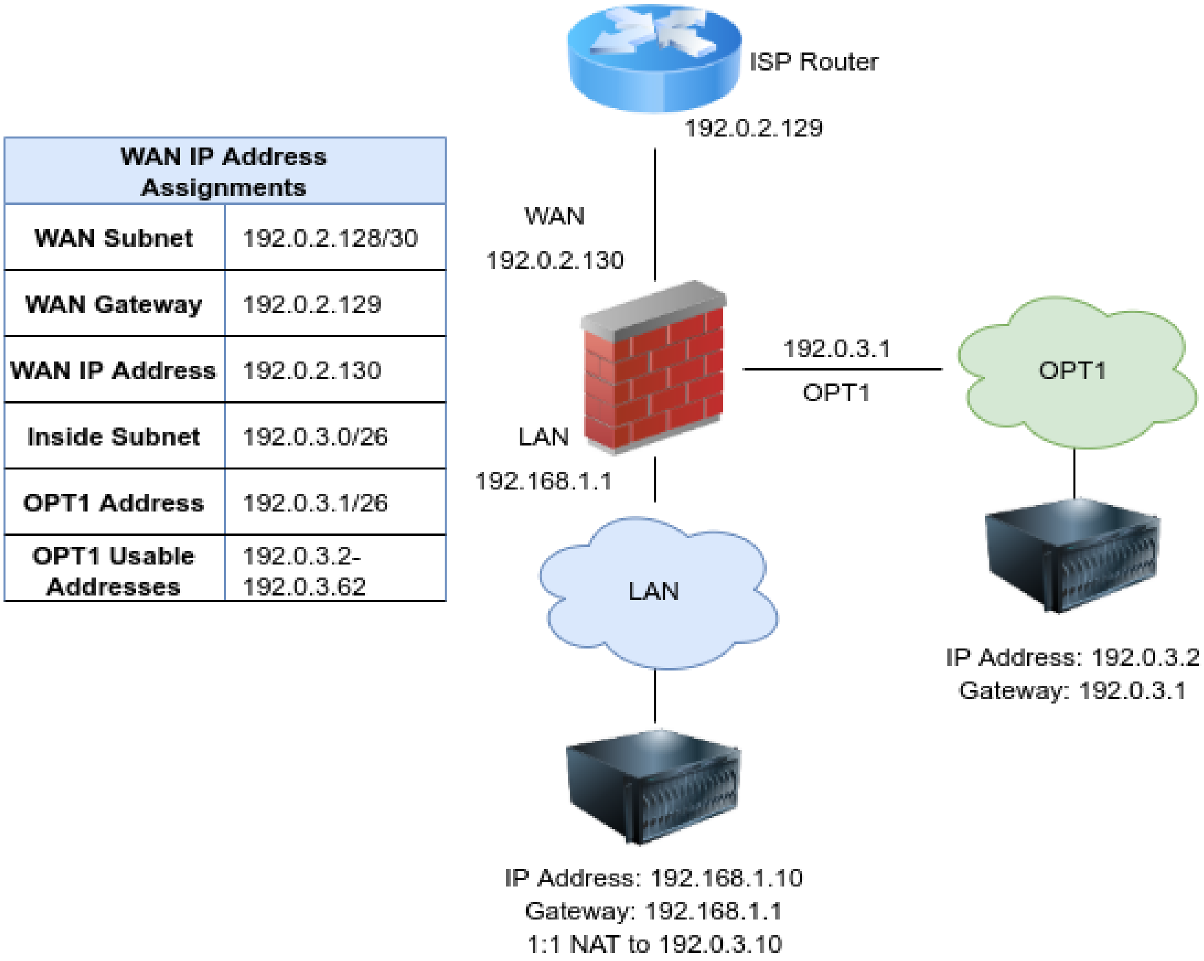

WAN IP Address Assignments	
WAN Subnet	192.0.2.128/30
WAN Gateway	192.0.2.129
WAN IP Address	192.0.2.130
Inside Subnet	192.0.3.0/26
OPT1 Address	192.0.3.1/26
OPT1 Usable Addresses	192.0.3.2- 192.0.3.62

Fig. 19: Multiple Public IP Addresses Using Two IP Subnets **Multiple IP subnets**

In other cases, a site may be allocated multiple IP subnets from the ISP. Usually when this happens, the site started with one of the two previously described arrangements, and later when requesting additional IP addresses the site was provided with an additional IP subnet. Ideally, this additional subnet will be routed to the firewall by the ISP, either to its WAN IP address in the case of a single firewall, or to a CARP VIP when using HA. If the provider refuses to route the IP subnet to the firewall, but rather routes it to their router and uses one of the IP addresses from the subnet as a gateway IP address, the firewall will need to use Proxy ARP VIPs, IP Alias VIPs, or a combination of IP Alias and CARP VIPs for the additional subnet. If at all possible, the provider should route the IP subnet to the firewall as it makes it easier to work with regardless of the firewall being used. It also eliminates the need to burn 3 IP addresses in the additional subnet, one for the network and broadcast addresses and one for the gateway IP address. With a routed subnet, the entire subnet is usable in combination with NAT.

Where the IP subnet is routed to the firewall, the scenario described in *Small WAN IP Subnet with Larger LAN IP Subnet* applies for an additional internal subnet. The subnet can be assigned to a new OPT interface, used it with NAT, or a combination of the two.

Additional IP Addresses via DHCP

Some ISPs require additional IP addresses to be obtained via DHCP. This is not a good means of obtaining multiple public IP addresses, and must be avoided in any serious network. A business-class connection should not require this. AZTCO-FW is one of the few firewalls which can be used in any capacity with additional IP addresses from DHCP. This offers limited flexibility in what the firewall can do with these addresses, leaving only two feasible options.

Bridging

If the additional IP addresses from DHCP must be directly assigned to the systems that will use them, bridging is the only option. Use an OPT interface bridged with WAN for these systems, and the systems must be configured to obtain their addresses using DHCP.

Pseudo multi-WAN

The only option for having the firewall pull these DHCP addresses as leases is a pseudo multi-WAN deployment. Install one network interface per public IP address, and configure each for DHCP. Plug all the interfaces into a switch between the firewall and the modem or router. Since the firewall will have multiple interfaces sharing a single broadcast domain, enable Suppress ARP messages on System > Advanced, Networking tab to eliminate ARP warnings in the system log, which are normal in this type of deployment.

The only use of multiple public IP addresses assigned in this fashion is for port forwarding. Port forwards can be used on each WAN interface that uses an IP address assigned to that interface by the ISP DHCP server. Outbound NAT to the OPT WANs will not work because of the limitation that each WAN must have a unique gateway IP address to properly direct traffic out of that WAN. This is discussed further in *Multiple WAN Connections*.

Virtual IP Addresses

AZTCO-FW® software enables the use of multiple IP addresses in conjunction with NAT or local services through Virtual IPs (VIPs).

There are four types of Virtual IP addresses available in AZTCO-FW: *IP Alias*, *CARP*, *Proxy ARP*, and *Other*. Each is useful in different situations. In most circumstances, AZTCO-FW will need to answer ARP request for a VIP which means that IP Alias, Proxy ARP or CARP must be used. In situations where ARP is not required, such as when additional public IP addresses are routed by a service provider to the WAN IP address on the firewall, use Other type VIPs.

AZTCO-FW will not respond to pings destined to Proxy ARP and Other type VIPs regardless of firewall rule configuration. With Proxy ARP and Other VIPs, NAT must be present on the firewall, forwarding traffic to an internal host for ping to function. See *Network Address Translation* for more information.

IP Alias

IP Aliases work like any other IP address on an interface, such as the actual interface IP address. They will respond to layer 2 (ARP) and can used as binding addresses by services on the firewall. They can also be used to handle multiple subnets on the same interface. AZTCO-FW will respond to ping on an IP Alias, and services on the firewall that bind to all interfaces will also respond on IP Alias VIPs unless the VIP is used to forward those ports in to another device (e.g. 1:1 NAT).

IP Alias VIPs can use *Localhost* as their interface to bind services using IP addresses from a block of routed addresses without specifically assigning the IP addresses to an interface. This is primarily useful in HA with CARP scenarios so that IP addresses do not need to be consumed by a CARP setup (one IP each per node, then the rest as CARP VIPs) when the subnet exists only inside the firewall (e.g. NAT or firewall services such as VPNs).

IP Aliases on their own do not synchronize to XMLRPC Configuration Synchronization peers because that would result in an IP address conflict. One exception to this is IP Alias VIPs using a CARP VIP "interface" for their interface. Those do not result in a conflict so they will synchronize. Another exception is IP Alias VIPs bound to Localhost as their interface. Because these are not active outside of the firewall itself, there is no chance of a conflict so they will also synchronize.

CARP

CARP VIPs are primarily used with High Availability redundant deployments utilizing CARP. CARP VIPs each have their own unique MAC address derived from their VHID, which can be useful even outside of a High Availability deployment.

See also:

For information on using CARP VIPs, see *High Availability*.

CARP VIPs may also be used with a single firewall. This is typically done in cases where the AZTCO-FW deployment will eventually be converted into an HA cluster node, or when having a unique MAC address is a requirement. In rare cases a provider requires each unique IP address on a WAN segment to have a distinct MAC address, which CARP VIPs provide.

CARP VIPs and IP Alias VIPs can be combined in two ways:

- To reduce the amount of CARP heartbeats by stacking IP Alias VIPs on CARP VIPs. See *Using IP Aliases to Reduce Heartbeat Traffic*.
- To use CARP VIPs in multiple subnets on a single interface. See *High Availability*.

Proxy ARP

Proxy ARP VIPs function strictly at layer 2, providing ARP replies for the specified IP address or CIDR range of IP addresses. This allows AZTCO-FW to accept traffic targeted at those addresses inside a shared subnet. For example, AZTCO-FW can forward traffic sent to an additional address inside its WAN subnet according to its NAT configuration. The address or range of addresses are not assigned to any interface on AZTCO-FW, because they don't need to be. This means no services on AZTCO-FW itself can respond on these IP addresses.

Proxy ARP VIPs do not sync to XML-RPC Configuration Sync peers because doing so would cause an IP address conflict.

Other

Other type VIPs define additional IP addresses for use when ARP replies for the IP address are not required. The only function of adding an *Other* type VIP is making that address available in the NAT configuration drop-down selectors. This is convenient when the firewall has a public IP block routed to its WAN IP address, IP Alias, or a CARP VIP.

Feature Comparison

Virtual IP Address Feature Comparison

This document summarizes and compares capabilities of the different Virtual IP Address types.

See *Virtual IP Addresses* for detailed information about each type of VIP.

VIP Features Table

Table 2: Virtual IP Address Feature Comparison

VIP Type	NAT	Binding	ARP/L2	Clustering	Subnet Mask	ICMP	Single/Range
IP Alias	Yes	Yes	Yes	See Notes	See Notes	Yes	Single
CARP	Yes	Yes	Yes	Yes	Yes	Yes	Single
Proxy ARP	Yes	No	Yes	No	n/a	No (1)	Either
Other	Yes	No	No	Yes (2)	n/a	No (1)	Either

Notes:

1. The ICMP Column represents responses from the firewall itself without NAT. With 1:1 NAT or port forwards, any VIP will pass ICMP through to the target device.
2. "Other" type VIPs are for routed subnets, and CARP is irrelevant, so they are compatible with HA (See below) **Virtual IP Feature Summary**

It is difficult to express all details of VIP capabilities in a table format, so this section contains a more thorough overview of the various types and what they can/cannot do a bullet point format.

IP Alias

- Can be used for NAT.
- Can be used by the firewall itself to bind/run services.
- Adds extra IP addresses to an interface.
- Generates ARP (Layer 2) responses for the VIP address.
- Can be in a different subnet than the real interface IP address when used directly on an interface.
- Will respond to ICMP ping if allowed by firewall rules.
- Must be added individually
- Subnet mask should match the interface IP, or /32. Matching the interface subnet is best. For IP addresses in different subnets at least one IP alias VIP must have the correct mask for the new subnet.
- Can be stacked on top of a CARP VIP to bypass VHID limits and lower the amount of CARP heartbeat traffic.
  - Stacked IP Alias VIPs will synchronize via XMLRPC.
  - Stacked IP Alias VIPs must be inside the same subnet as the CARP VIP upon which they are placed.
- Can be added to localhost for binding services in routed subnets. IP Alias VIPs bound to localhost will synchronize via XMLRPC

CARP

- Can be used for NAT.
- Can be used by the firewall itself to bind/run services.
- Generates ARP (Layer 2) traffic for the VIP.

- Can be used for clustering (master firewall and standby failover firewall.)
- CARP VIPs may be in other subnets.
- Will respond to ICMP ping if allowed by firewall rules.
- Must be added individually.
- Subnet mask must match the interface IP address.
- Generates its own MAC address for the VIP. This MAC is different than its physical parent interface.

Proxy ARP

- Can be used for NAT.
- Cannot be used by the firewall itself to bind/run services.
- Generates ARP (Layer 2) traffic for the VIP.
- Can be in a different subnet than the real interface IP.
- Will not respond to ICMP ping.
- Can be added individually or as a subnet to make a group of VIPs.

Other

- Can be used for NAT.
- Cannot be used by the firewall itself to bind/run services.
- Can be used if the address is routed to the firewall without needing ARP/Layer 2 messages. (e.g. Upstream provider routes a subnet to the WAN IP address)
- Can be in a different subnet than the real interface IP address.
- Will not respond to ICMP echo requests.
- Can be added individually or as a subnet to make a group of VIPs.
- Can be used with CARP, e.g. subnet routed to external CARP VIP.

Using EasyRule to Add Firewall Rules

The EasyRule function found in the GUI and on the command line can add firewall rules quickly.

EasyRule in the GUI

In the AZTCO-FW® software GUI, this function is available in the Firewall Log view (Status > System Logs, Firewall tab).

The ⊟ icon next to the source IP address adds a block rule for that IP address on the interface. To be more precise, it creates or adds to an alias containing IP addresses added from Easy Rule and blocks them on the selected interface.

The ⊞ icon next to the destination IP address works similar to the block action, but it adds a more precise pass rule. This pass rule allows traffic on the interface but it must match the same protocol, source IP address, destination IP address, and destination port.

EasyRule in the Shell

The shell version of Easy Rule, easyrule, can add a firewall rule from a shell prompt. When the easyrule command is run without parameters, it prints a usage message to explain its syntax.

The way easyrule adds a block rule using an alias, or a precise pass rule specifying the protocol, source, and destination, work similar to the GUI version.

```
: easyrule usage:
Blocking only requires an IP to block easyrule block <interface>
        <source IP>

Passing requires more detail, as it must be as specific as possible. The destination␣
↪port is optional if the protocol does not require a port (e.g. ICMP, OSPF, etc).
        easyrule pass <interface> <protocol> <source IP> <destination ip> [destination ␣↪port]

Block example:
        easyrule block wan 1.2.3.4

Pass example (protocol with port): easyrule pass wan tcp 1.2.3.4
        192.168.0.4 80

Pass example (protocol without port):
        easyrule pass wan icmp 1.2.3.4 192.168.0.4
```

The source code of those scripts can be adapted for adding firewall rules in other ways, but that is left as an exercise for the reader.

See also:

- *Ordering of NAT and Firewall Processing*
- *Viewing the Firewall Log*
- *Filter Reload Status*

8.1.2 Aliases

Aliases are collections of addresses that allow many hosts to be acted upon by a small number of firewall rules. They can greatly simplify a ruleset and make it easier to understand and manage.

Aliases

Aliases define a group ports, hosts, or networks. Aliases can be referenced by firewall rules, port forwards, outbound NAT rules, and other places in the firewall GUI. Using aliases results in significantly shorter, self-documenting, and more manageable rulesets.

Note: Do not confuse Aliases in this context with interface IP aliases, which are a means of adding additional IP addresses to a network interface.

Alias Basics

Aliases are located at Firewall > Aliases. The page is divided into separate tabs for each type of alias: IP, Ports, URLs, and the All tab which shows every alias in one large list. When creating an alias, add it to any tab and it will be sorted to the correct location based on the type chosen.

The following types of aliases can be created:

Host Aliases containing single IP addresses or hostnames

Network Aliases containing CIDR-masked lists of networks, hostnames, IP address ranges, or single IP addresses

Port These aliases contain lists of port numbers or ranges of ports for TCP or UDP.

URL The alias is built from the file at the specified URL but is read only a single time, and then becomes a normal network or port type alias.

URL Table The alias is built from the file at the specified URL but is updated by fetching the list from the URL periodically.

Each alias type is described in more detail throughout this section.

Nesting Aliases

Most aliases can be nested inside of other aliases so long as they are the same type. For example, one alias can nest an alias containing web servers, an alias containing mail servers, and a servers alias that contains both the web and mail server aliases all together in one larger Servers alias. URL Table aliases cannot be nested.

Using Hostnames in Aliases

Hostnames can also be used in aliases. Any hostname can be entered into a host or network alias and it will be periodically resolved and updated by the firewall. If a hostname returns multiple IP addresses, all of the returned IP addresses are added to the alias. This is useful for tracking dynamic DNS entries to allow specific users into services from dynamic IP addresses.

Note: This feature is *not* useful for allowing or disallowing users to large public web sites. Large and busy sites tend to have constantly rotating or random responses to DNS queries so the contents of the alias do not necessarily match up with the response a user will receive when they attempt to the resolve the same site name. It can work for smaller sites that have only a few servers and do not include incomplete sets of addresses in their DNS responses.

Mixing IPv4 and IPv6 Addresses in Aliases

IPv4 and IPv6 addresses can be mixed inside an alias. The firewall will use the appropriate type of addresses when the alias is referenced in a specific rule.

Alias Sizing Concerns

The total size of all tables must fit in roughly half the amount of Firewall Maximum Table Entries, which defaults to 200,000. If the maximum number of table entries is not large enough to contain all of the entries, the rules may fail to load. See *Firewall Maximum Table Entries* for information on changing that value. The aliases must fit in twice in the total area because of the way aliases are loaded and reloaded; The new list is loaded alongside the old list and then the old one is removed.

This value can be increased as much required, provided that the firewall contains sufficient RAM to hold the entries. The RAM usage is similar to, but less than, the state table but it is still safe to assume 1K per entry to be conservative.

Configuring Aliases

To add an alias:

- Navigate to Firewall > Aliases
- Click + Add
- Enter a Name for the alias. The name may only consist of the characters a-z, A-Z, 0-9 and _.
- Enter a Description for the alias itself
- Select the Type for the alias. The various types are discussed throughout this section.
- Enter the type-specific information as needed. Each type has an data field and a description field for each entry.

To add new members to an alias, + clickAdd at the bottom of the list of entries.

To remove members from an alias, 🗑 clickDelete at the end of the row to remove.

When the alias is complete, click Save to store the alias contents.

Each manually entered alias is limited to 5,000 members, but some browsers have trouble displaying or using the page with more than around 3,000 entries. For large numbers of entries, use a *URL Table* type alias which is capable of handling larger lists.

Host Aliases

Host type aliases contain groups of IP addresses. Figure *Example Hosts Alias* shows an example of a host type alias used to contain a list of public web servers.

Other host type aliases can be nested inside this entry. Hostnames may also be used as entries, as explained previously.

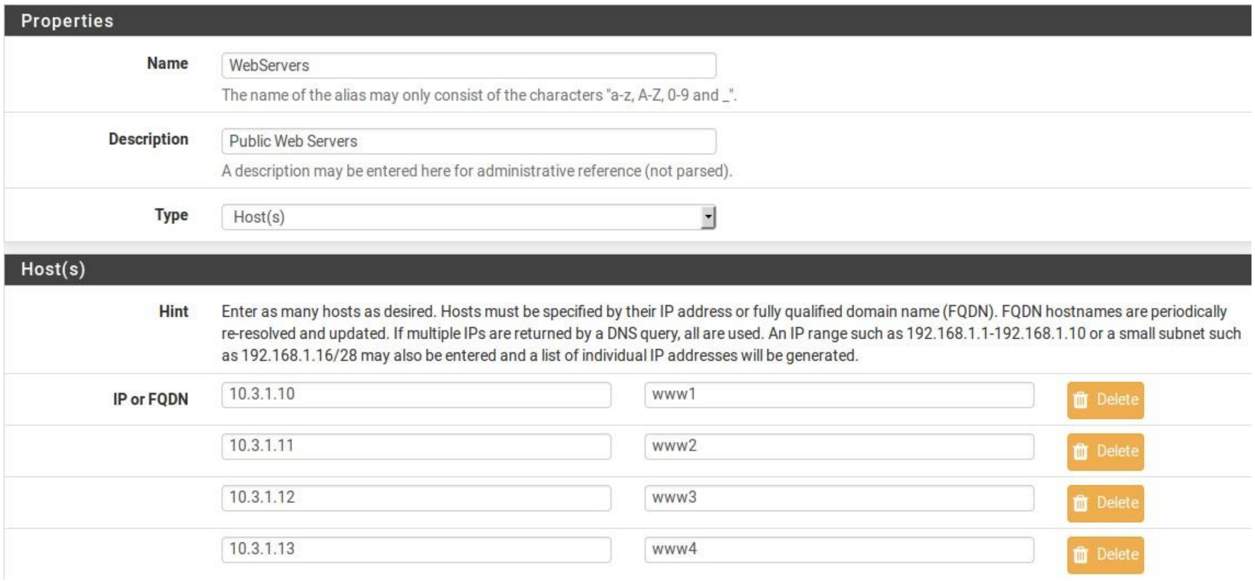

Properties

Name	WebServers The name of the alias may only consist of the characters "a-z, A-Z, 0-9 and _".		
Description	Public Web Servers A description may be entered here for administrative reference (not parsed).		
Type	Host(s)		

Host(s)

Hint	Enter as many hosts as desired. Hosts must be specified by their IP address or fully qualified domain name (FQDN). FQDN hostnames are periodically re-resolved and updated. If multiple IPs are returned by a DNS query, all are used. An IP range such as 192.168.1.1-192.168.1.10 or a small subnet such as 192.168.1.16/28 may also be entered and a list of individual IP addresses will be generated.		
IP or FQDN	10.3.1.10	www1	Delete
	10.3.1.11	www2	Delete
	10.3.1.12	www3	Delete
	10.3.1.13	www4	Delete

Fig. 20: Example Hosts Alias

Network Aliases

Network type aliases contain groups of networks or IP address ranges. Single hosts can also be included in network aliases by selecting a */32* network mask for IPv4 addresses or a */128* prefix length for IPv6 addresses. Figure *Example Network Alias* shows an example of a network alias that is used later in this chapter.

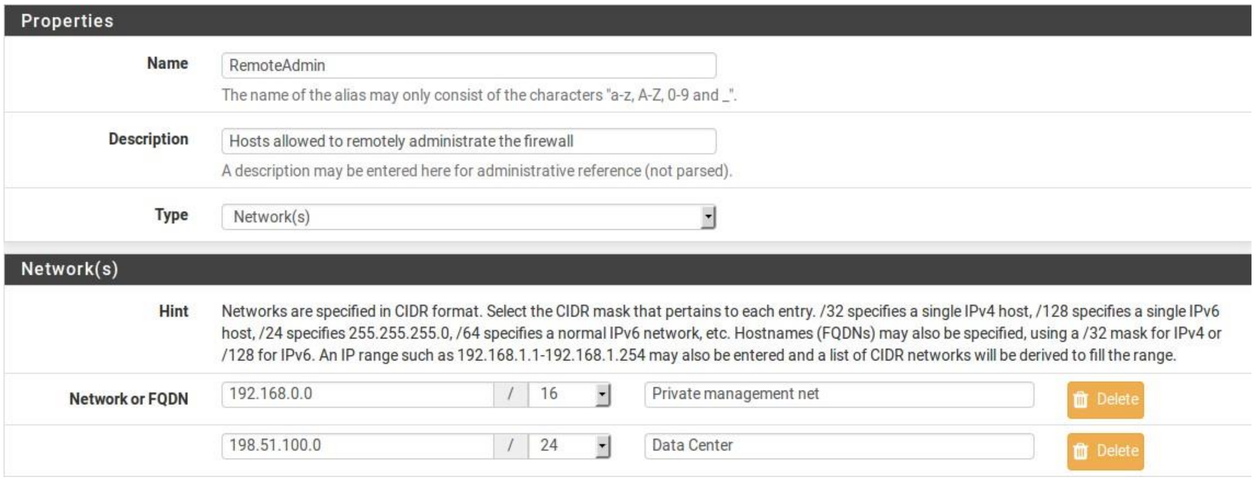

Properties

Name	RemoteAdmin The name of the alias may only consist of the characters "a-z, A-Z, 0-9 and _".		
Description	Hosts allowed to remotely administrate the firewall A description may be entered here for administrative reference (not parsed).		
Type	Network(s)		

Network(s)

Hint	Networks are specified in CIDR format. Select the CIDR mask that pertains to each entry. /32 specifies a single IPv4 host, /128 specifies a single IPv6 host, /24 specifies 255.255.255.0, /64 specifies a normal IPv6 network, etc. Hostnames (FQDNs) may also be specified, using a /32 mask for IPv4 or /128 for IPv6. An IP range such as 192.168.1.1-192.168.1.254 may also be entered and a list of CIDR networks will be derived to fill the range.		
Network or FQDN	192.168.0.0 / 16	Private management net	Delete
	198.51.100.0 / 24	Data Center	Delete

Fig. 21: Example Network Alias

Other host or network aliases can be nested inside this entry. Hostnames may also be used as entries, as explained previously.

When an alias entry contains an IPv4 range it is automatically translated by the firewall to an equivalent set of IPv4 CIDR networks that will exactly contain the provided range. As shown in Figure *Example IP Range After*, the range is expanded when the alias is saved, and the resulting list of IPv4 CIDR networks will match exactly the requested range, nothing more, nothing less.

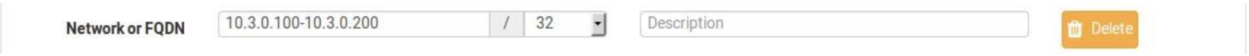

Network or FQDN	10.3.0.100-10.3.0.200	/	32	Description	Delete

Fig. 22: Example IP Range Before

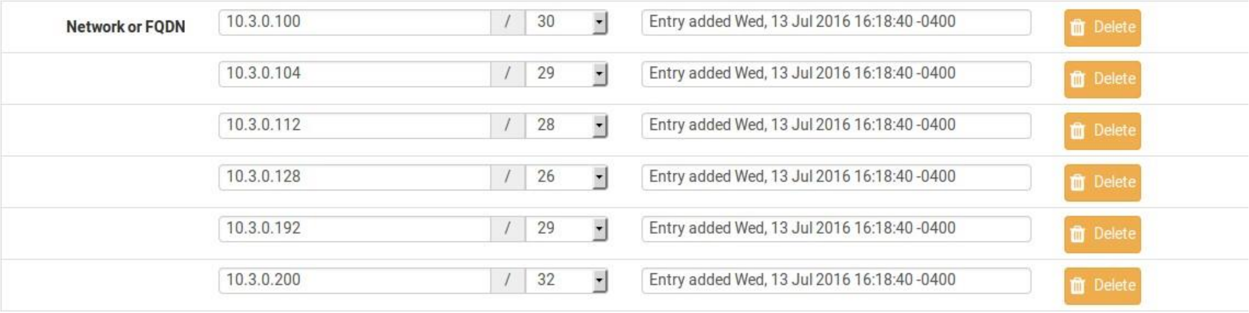

Network or FQDN	10.3.0.100	/	30	Entry added Wed, 13 Jul 2016 16:18:40 -0400	Delete
	10.3.0.104	/	29	Entry added Wed, 13 Jul 2016 16:18:40 -0400	Delete
	10.3.0.112	/	28	Entry added Wed, 13 Jul 2016 16:18:40 -0400	Delete
	10.3.0.128	/	26	Entry added Wed, 13 Jul 2016 16:18:40 -0400	Delete
	10.3.0.192	/	29	Entry added Wed, 13 Jul 2016 16:18:40 -0400	Delete
	10.3.0.200	/	32	Entry added Wed, 13 Jul 2016 16:18:40 -0400	Delete

Fig. 23: Example IP Range After

Port Aliases

Port type aliases contain groups of ports and port ranges. The protocol is not specified in the alias; The firewall rule where the alias is used will define the protocol as TCP, UDP, or both. Figure *Example Ports Alias* shows an example of a port type alias.

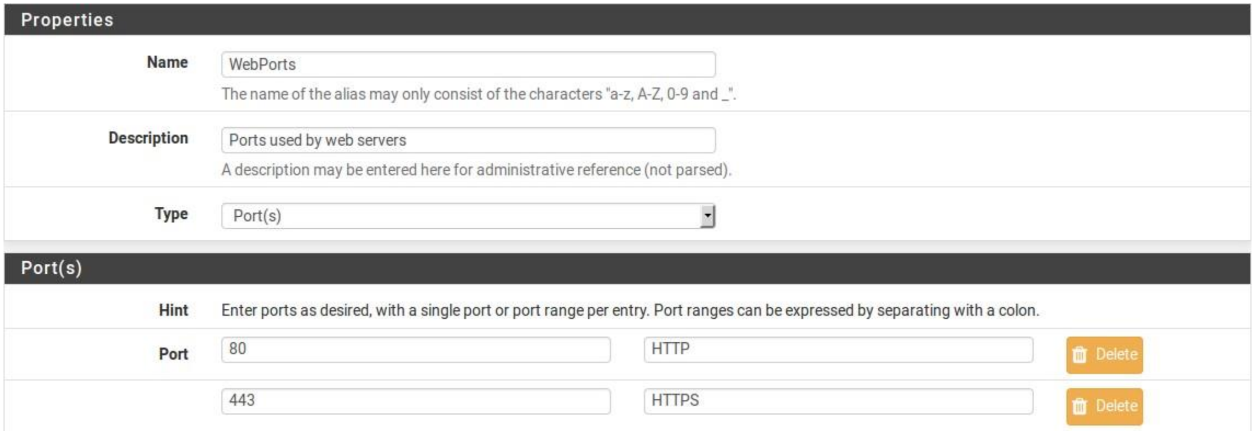

Properties

Name: WebPorts
The name of the alias may only consist of the characters "a-z, A-Z, 0-9 and _".

Description: Ports used by web servers
A description may be entered here for administrative reference (not parsed).

Type: Port(s)

Port(s)

Hint: Enter ports as desired, with a single port or port range per entry. Port ranges can be expressed by separating with a colon.

Port	80	HTTP	Delete
	443	HTTPS	Delete

Fig. 24: Example Ports Alias

Enter another port-type alias name into the Port field to nest other port- type aliases inside this alias.

URL Aliases

With a URL type alias, a URL is set which points to a text file that contains a list of entries. Multiple URLs may be entered. When Save is clicked, up to 3,000 entries from each URL are read from the file and imported into a network type alias.

If *URL (IPs)* is selected, then the URLs must contain IP address or CIDR masked network entries, and the firewall creates a network type alias from the contents.

If *URL (Ports)* is selected, then the URL must contain only port numbers or ranges, and the firewall creates a port type alias from the contents.

URL Table Aliases

A URL Table alias behaves in a significantly different way than the URL alias. For starters, it does not import the contents of the file into a normal alias. It downloads the contents of the file into a special location on the firewall and uses the contents for what is called a persist table, also known as a file-based alias. The full contents of the alias are not directly editable in the GUI, but can be viewed in the Tables viewer (See *Viewing the Contents of Tables*).

For a URL Table alias, the drop-down list after the / controls how many days must pass before the contents of the alias are re-fetched from the stored URL by the firewall. When the time comes, the alias contents will be updated overnight by a script which re-fetches the data.

URL Table aliases can be quite large, containing many thousands of entries. Some customers use them to hold lists of all IP blocks in a given country or region, which can easily surpass 40,000 entries. The pfBlocker package uses this type of alias when handling country lists and other similar actions.

Currently, URL Table aliases are not capable of being nested.

If *URL Table (IPs)* is selected, then the URLs must contain IP address or CIDR masked network entries, and the firewall creates a network type alias from the contents.

If *URL Table (Ports)* is selected, then the URL must contain only port numbers or ranges, and the firewall creates a port type alias from the contents.

Bulk Import Network Aliases

Another method of importing multiple entries into an alias is to use the bulk import feature.

To use the import feature:

- Navigate to Firewall > Aliases
- Click Import
- Fill in the Alias Name and Description
- Enter the alias contents into the Aliases to import text area, one entry per line.
- Click Save

Common usage examples for this page include lists of IP addresses, networks, and blacklists. The list may contain IP addresses, CIDR masked networks, IP ranges, or port numbers. The firewall will attempt to determine the target alias type automatically.

The firewall imports items into a normal alias which can be edited later.

Using Aliases

When a letter is typed into an input box which supports aliases, a list of matching aliases is displayed. Select the desired alias from the list, or type its name out completely.

Note: Alias autocompletion is not case sensitive but it is restricted by type. For example, a Network or Host type alias will be listed in autocomplete for a Network field, but a Port alias will not; A port alias can be used in a port field, but a Network alias will not be in the list.

Figure *Autocompletion of Hosts Alias* shows how the WebServers alias, configured as shown in Figure *Example Hosts Alias*, can be used in the Destination field when adding or editing a firewall rule.

- Edit the firewall rule
- Select *Single host or alias*
- Then type the first letter of the desired alias: Enter W and the alias appears as shown.

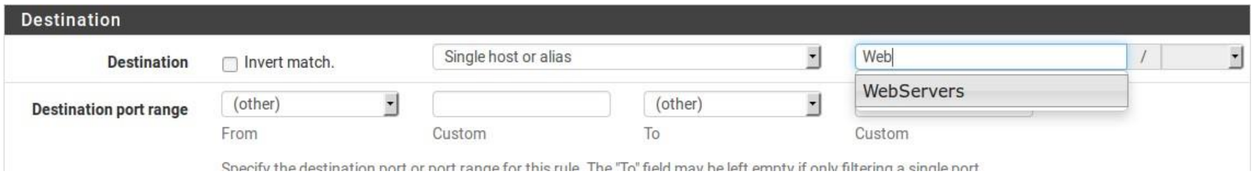

Destination

Destination □ Invert match. Single host or alias Web /

WebServers

Destination port range (other) From Custom (other) To Custom

Specify the destination port or port range for this rule. The "To" field may be left empty if only filtering a single port

Fig. 25: Autocompletion of Hosts Alias

Figure *Autocompletion of Ports Alias* shows the autocompletion of the ports alias configured as shown in Figure *Example Ports Alias*. If multiple aliases match the letter entered, all matching aliases of the appropriate type are listed. Click on the desired alias to select it.

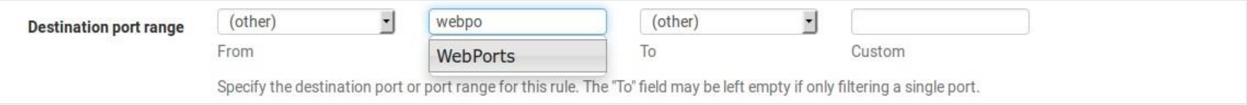

Destination port range (other) From webpo

WebPorts

(other) To Custom

Specify the destination port or port range for this rule. The "To" field may be left empty if only filtering a single port.

Fig. 26: Autocompletion of Ports Alias

Figure *Example Rule Using Aliases* shows the rule created using the WebServers and WebPorts aliases. This rule is on WAN, and allows any source to the IP addresses defined in the WebServers alias when using the ports defined in the WebPorts alias.

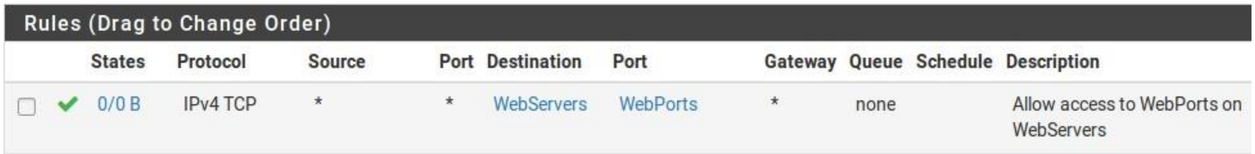

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description
□	✓	0/0 B	IPv4 TCP	*	*	WebServers	WebPorts	*	none		Allow access to WebPorts on WebServers

Fig. 27: Example Rule Using Aliases

Hovering the mouse cursor over an alias on the Firewall > Rules page shows a tooltip displaying the contents of the alias with the descriptions included in the alias. Figure *Hovering Shows Hosts Contents* shows this for the WebServers alias and Figure *Hovering Shows Ports Contents* for the ports alias.

8.1.3 Firewall Guides

How to perform various tasks with firewall rules.

See also:

- *Blocking Web Sites*
- *Allowing Remote Access to the GUI*
- *Preventing RFC1918 Traffic from Exiting a WAN Interface*
- *Configuring AZTCO-FW software for Online Gaming*

Troubleshooting problems with firewall behavior.

See also:

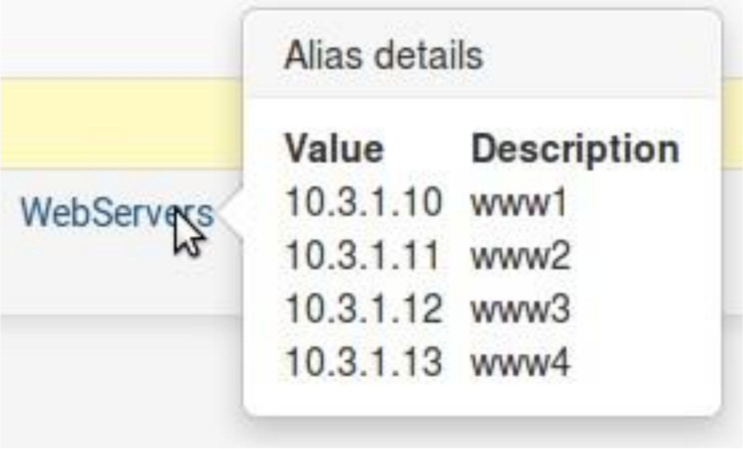

Alias details

Value	Description
10.3.1.10	www1
10.3.1.11	www2
10.3.1.12	www3
10.3.1.13	www4

Fig. 28: Hovering Shows Hosts Contents

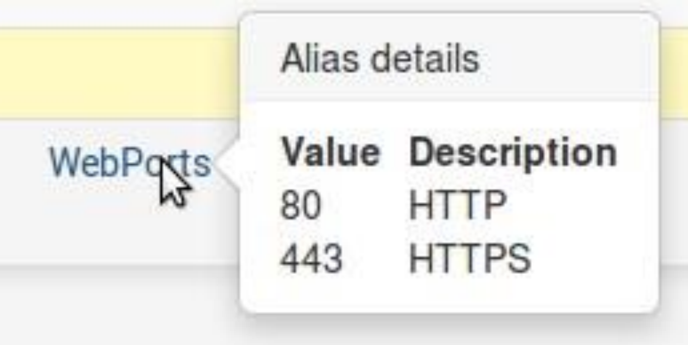

Alias details

Value	Description
80	HTTP
443	HTTPS

Fig. 29: Hovering Shows Ports Contents

- *Troubleshooting Thread Errors with Hostnames in Aliases*
- *Troubleshooting Firewall Rules*
- *Troubleshooting Asymmetric Routing*
- *Troubleshooting Blocked Log Entries for Legitimate Connection Packets*

CHAPTER

NINE

NETWORK ADDRESS TRANSLATION

9.1 Port Forwards

Port forwards allow access to a specific port, port range or protocol on a privately addressed internal network device. The name "port forward" was chosen because it is what most people understand in this context, and it was renamed from the more technically appropriate "Inbound NAT" after countless complaints from confused users. Similar functionality is also called "Destination NAT" in other products. However, "Port Forward" a misnomer, as port forward rules can redirect GRE and ESP protocols in addition to TCP and UDP ports, and it can be used for various types of traffic redirection as well as traditional port forwards. This is most commonly used when hosting servers, or using applications that require inbound connections from the Internet.

9.1.1 Risks of Port Forwarding

In a default configuration, AZTCO-FW® does not let in any traffic initiated from hosts on the Internet. This provides protection from anyone scanning the Internet looking for systems to attack. When a port forward rule exists, AZTCO-FW will allow any traffic matching the corresponding firewall rule. It does not know the difference between a packet with a malicious payload and one that is benign. If the connection matches the firewall rule, it is allowed. Host based controls must be used by the target system to secure any services allowed through the firewall.

9.1.2 Port Forwarding and Local Services

Port forwards take precedence over any services running locally on the firewall, such as the web interface, SSH, and so on. For example this means if remote web interface access is allowed from the WAN using HTTPS on TCP port 443, a port forward on WAN for TCP 443 will take precedence and the web interface will no longer be accessible from WAN. This does not affect access on other interfaces, only the interface containing the port forward.

9.1.3 Port Forwarding and 1:1 NAT

Port forwards also take precedence over 1:1 NAT. If a port forward is defined on one external IP address forwarding a port to a host, and a 1:1 NAT entry is also defined on the same external IP address forwarding everything into a different host, then the port forward remains active and continues forwarding to the original host.

9.1.4 Adding Port Forwards

Port Forwards are managed at Firewall > NAT, on the Port Forward tab. The rules on this screen are managed in the same manner as firewall rules (see *Introduction to the Firewall Rules screen*).

To begin adding a port forward entry, click ↧ Add button to reach the Port Forward editing screen. The following options are available for port forwards:

Disable A checkbox to optionally Disable this NAT port forward. To deactivate the rule, check this box.

No RDR (NOT) Negates the meaning of this port forward, indicating that no redirection should be performed if this rule is matched. Most configurations will not use this field. This would be used to override a forwarding action, which may be needed in some cases to allow access to a service on the firewall on an IP being used for 1:1 NAT, or another similar advanced scenario.

Interface The interface where the port forward will be active. In most cases this will be WAN. For additional WAN links or local redirects this may be different interface. The Interface is the location on the firewall where traffic for this port forward enters.

Protocol The Protocol of the incoming traffic to match. This must be set to match the type of service being forwarded, whether it is *TCP*, *UDP*, or another available choice. Most common services being forwarded will be *TCP* or *UDP*, but consult the documentation for the service or even a quick web search to confirm the answer. The *TCP/UDP* option forwards both TCP and UDP together in a single rule.

Source These options are hidden behind an Advanced button by default, and set to *any* source. The Source options restrict which source IP addresses and ports can access this port forward entry.

These are not typically necessary. If the port forward must be reachable from any location on the Internet, the source must be *any*. For restricted access services, use an alias here so only a limited set of IP addresses may access the port forward. Unless the service absolutely requires a specific source port, the Source Port Range must be left as *any* since nearly all clients will use randomized source ports.

Destination The IP address where the traffic to be forwarded is initially destined. For port forwards on WAN, in most cases this is *WAN Address*. Where multiple public IP addresses are available, it may be a Virtual IP (see *Virtual IP Addresses*) on WAN.

Destination port range The original destination port of the traffic, as it is coming in from the Internet, before it is redirected to the specified target host. If forwarding a single port, enter it in the From port box and leave the To port box blank. A list of common services is available to choose from in the drop down boxes in this group. Port aliases may also be used here to forward a set of services. If an alias is used here, the same alias must be used as the Redirect target port.

Redirect target IP The IP address where traffic will be forwarded, or technically redirected. An alias here, but the alias must only contain a single address. If the alias contains multiple addresses, the port will be forwarded to each host alternately, which is not what most people want. To setup load balancing for one port to multiple internal servers, see *HAProxy package*.

Redirect target port Where the forwarded port range will begin. If a range of ports is forwarded, e.g. 19000-19100, only the local starting point is specified since the number of ports must match up one-to-one.

This field allows opening a different port on the outside than the host on the inside is listening on. For example external port 8888 may forward to local port 80 for HTTP on an internal server. A list of common services is available to pick from in the drop down box.

Port aliases may also be used here to forward a set of services. If an alias is used here, the same alias must be used as the Destination port range.

Description As in other parts of AZTCO-FW, this field is available for a short sentence about what the port forward does or why it exists.

No XML-RPC Sync This option is only relevant if an HA Cluster configuration is in use, and should be skipped otherwise. When using an HA cluster with configuration synchronization, checking this box will prevent the rule from being synchronized to the other members of a cluster (see *High Availability*). Typically all rules should synchronize, however. This option is only effective on master nodes, it does *not* prevent a rule from being overwritten on slave nodes.

NAT Reflection This topic is covered in more detail later in this chapter (*NAT Reflection*). This option allows reflection to be enabled or disabled a per-rule basis to override the global default. The options in this field are explained in more detail in *NAT Reflection*.

Filter Rule Association This final option is *very* important. A port forward entry only defines which traffic will be redirected, a firewall rule is required to *pass* any traffic through that redirection. By default, *Add associated filter rule* is selected. The available choices are:

None If this is chosen, no firewall rule will be created.

Add associated filter rule This option creates a firewall rule that is linked to this NAT port forward rule. Changes made to the NAT rule are updated in the firewall rule automatically. This is the best choice for most use cases. If this option is chosen, after the rule is saved a link is placed here which leads to the associated firewall rule.

Add unassociated filter rule This option creates a firewall rule that separate from this NAT port forward. Changes made to the NAT rule must be manually changed in the firewall rule. This can be useful if other options or restrictions must be set on the firewall rule rather than the NAT rule.

Pass This choice uses a special pf keyword on the NAT port forward rule that causes traffic to be passed through without the need of a firewall rule. Because no separate firewall rule exists, any traffic matching this rule is forwarded in to the target system.

Note: Rules using *Pass* will only work on the interface containing the default gateway for the firewall, so they do not work effectively with Multi-WAN.

- Click Save
- Click Apply Changes

Figure *Port Forward Example* contains an example of the port forward editing screen filled in with the proper settings to forward HTTP inbound on WAN destined to the WAN IP address to the internal system at 10.3.0.15.

After clicking Save, the port forward list is displayed again, and the newly created entry will be present in the list, as in Figure *Port Forward List*.

Double check the firewall rule, as seen under Firewall > Rules on the tab for the interface upon which the port forward was created. The rule will show that traffic is allowed into the internal IP address on the proper port, as shown in Figure *Port Forward Firewall Rule*.

The *Source* of the automatically generated rule should be restricted where possible. For things such as mail and web servers that typically need to be widely accessible, this isn’t practical, but for remote management services such as SSH, RDP and others, there are likely only a small number of hosts that should be able to connect using those protocols

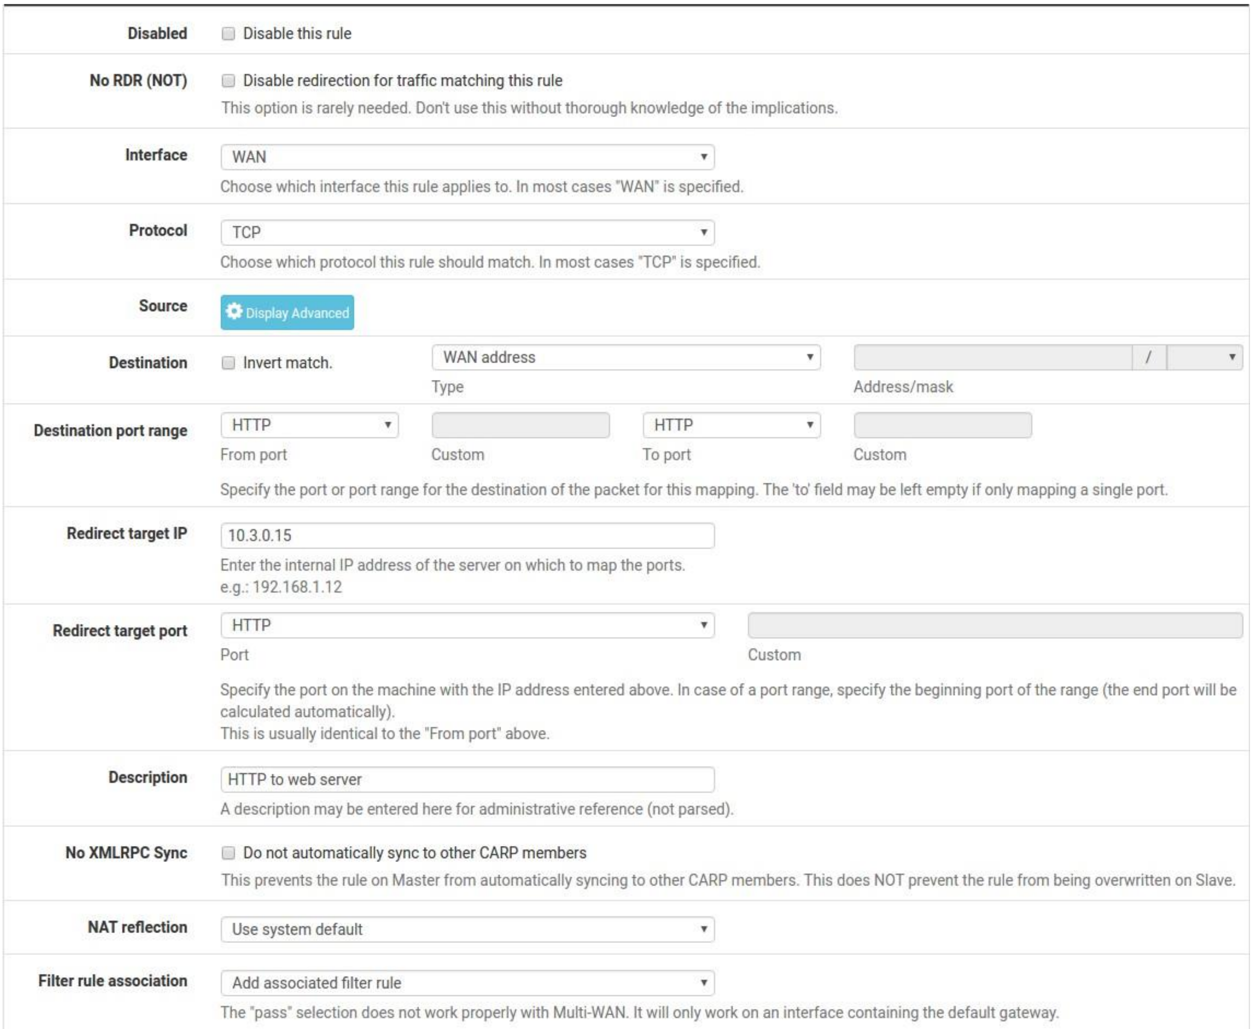

Disabled	□ Disable this rule
No RDR (NOT)	□ Disable redirection for traffic matching this rule This option is rarely needed. Don't use this without thorough knowledge of the implications.
Interface	WAN Choose which interface this rule applies to. In most cases "WAN" is specified.
Protocol	TCP Choose which protocol this rule should match. In most cases "TCP" is specified.
Source	Display Advanced
Destination	□ Invert match. WAN address (Type) / (Address/mask)
Destination port range	HTTP (From port) (Custom) HTTP (To port) (Custom) Specify the port or port range for the destination of the packet for this mapping. The 'to' field may be left empty if only mapping a single port.
Redirect target IP	10.3.0.15 Enter the internal IP address of the server on which to map the ports. e.g.: 192.168.1.12
Redirect target port	HTTP (Port) (Custom) Specify the port on the machine with the IP address entered above. In case of a port range, specify the beginning port of the range (the end port will be calculated automatically). This is usually identical to the "From port" above.
Description	HTTP to web server A description may be entered here for administrative reference (not parsed).
No XMLRPC Sync	□ Do not automatically sync to other CARP members This prevents the rule on Master from automatically syncing to other CARP members. This does NOT prevent the rule from being overwritten on Slave.
NAT reflection	Use system default
Filter rule association	Add associated filter rule The "pass" selection does not work properly with Multi-WAN. It will only work on an interface containing the default gateway.

Fig. 1: Port Forward Example

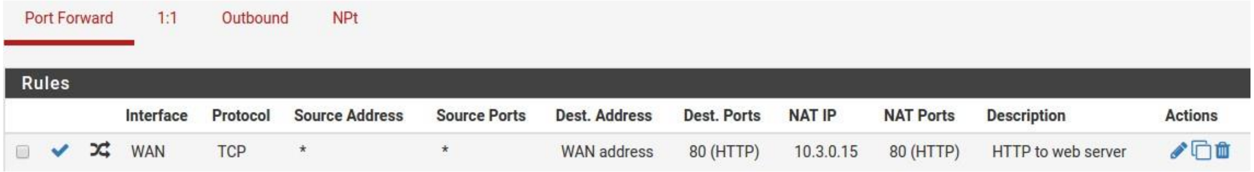

Port Forward 1:1 Outbound NPt

Rules

			Interface	Protocol	Source Address	Source Ports	Dest. Address	Dest. Ports	NAT IP	NAT Ports	Description	Actions
□	✓	⤮	WAN	TCP	*	*	WAN address	80 (HTTP)	10.3.0.15	80 (HTTP)	HTTP to web server	

Fig. 2: Port Forward List

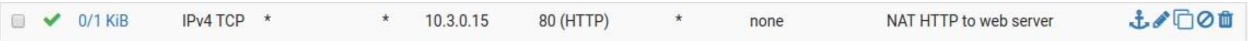

□	✓	0/1 KiB	IPv4 TCP	*	*	10.3.0.15	80 (HTTP)	*	none	NAT HTTP to web server	

Fig. 3: Port Forward Firewall Rule into a server from across the Internet. A much more secure practice is to create an alias of authorized hosts, and then change the source from *any* to the alias. Otherwise, the server is wide

open to the entire Internet. Test the port forward first with the unrestricted source, and after verifying it works, restrict the source as desired.

If everything looks right, the port forward will work when tested from outside the network. If something went wrong, see *Port Forward Troubleshooting* later in this chapter.

9.1.5 Tracking Changes to Port Forwards

As mentioned in Figure *Firewall Rule Time Stamps* for firewall rules, a timestamp is added to a port forward entry when it is created or last edited, to show which user created the rule, and the last person to edit the rule. Firewall rules automatically created by associated NAT rules are also marked as such on the associated firewall rule's creation timestamp.

9.1.6 Port Forward Limitations

A single port can only be forwarded to one internal host for each available public IP address. For instance, if only one public IP address is available, one internal web server that uses TCP port 80 to serve web traffic can be configured. Any additional servers must use alternate ports such as 8080. If five available public IP addresses are configured as Virtual IP addresses, then five internal web servers using port 80 can be configured. See *Virtual IP Addresses* for more about Virtual IP addresses.

There is one uncommon but sometimes applicable exception to this rule. If a particular port must be forwarded to a specific internal host only for certain source IP addresses, and that same port can be forwarded to a different host for other source IP addresses, that is possible by specifying the source address in the port forward entries, such as in Figure *Port Forward Example with Different Sources*.

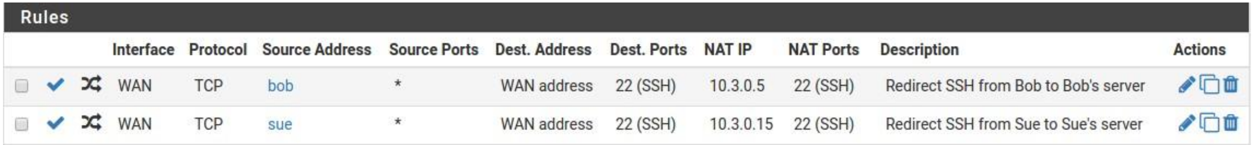

Rules

			Interface	Protocol	Source Address	Source Ports	Dest. Address	Dest. Ports	NAT IP	NAT Ports	Description	Actions
□	✓	⤮	WAN	TCP	bob	*	WAN address	22 (SSH)	10.3.0.5	22 (SSH)	Redirect SSH from Bob to Bob's server	
□	✓	⤮	WAN	TCP	sue	*	WAN address	22 (SSH)	10.3.0.15	22 (SSH)	Redirect SSH from Sue to Sue's server	

Fig. 4: Port Forward Example with Different Sources

In order for port forwards on WAN addresses to be accessible by using their respective WAN IP address from internalfacing interfaces, NAT reflection must be enabled, which is described in *NAT Reflection*. Always test port forwards from a system on a different Internet connection, and not from inside the network. Testing from a mobile device on 3G/4G is a quick and easy way to confirm external connectivity.

9.1.7 Service Self-Configuration With UPnP or NAT-PMP

Some programs support Universal Plug-and-Play (UPnP) or NAT Port Mapping Protocol (NAT-PMP) to automatically configure NAT port forwards and firewall rules. Even more security concerns apply there, but in home use the benefits often outweigh any potential concerns. See *UPnP & NAT-PMP* for more information on configuring and using UPnP and NAT-PMP.

9.1.8 Traffic Redirection with Port Forwards

Another use of port forwards is for transparently redirecting traffic from an internal network. Port forwards specifying the *LAN* interface or another internal interface will redirect traffic matching the forward to the specified destination. This is most commonly used for transparently proxying HTTP traffic to a proxy server, or redirecting all outbound DNS to one server.

The NAT entries shown in Figure *Example Redirect Port Forward* are an example of a configuration that will redirect all HTTP traffic coming into the LAN interface to Squid (port 3128) on the host 10.3.0.10, but will not redirect the traffic coming from the actual squid proxy itself. They must be in the correct order in the list of port forwards: The negate rule first, then the redirect.

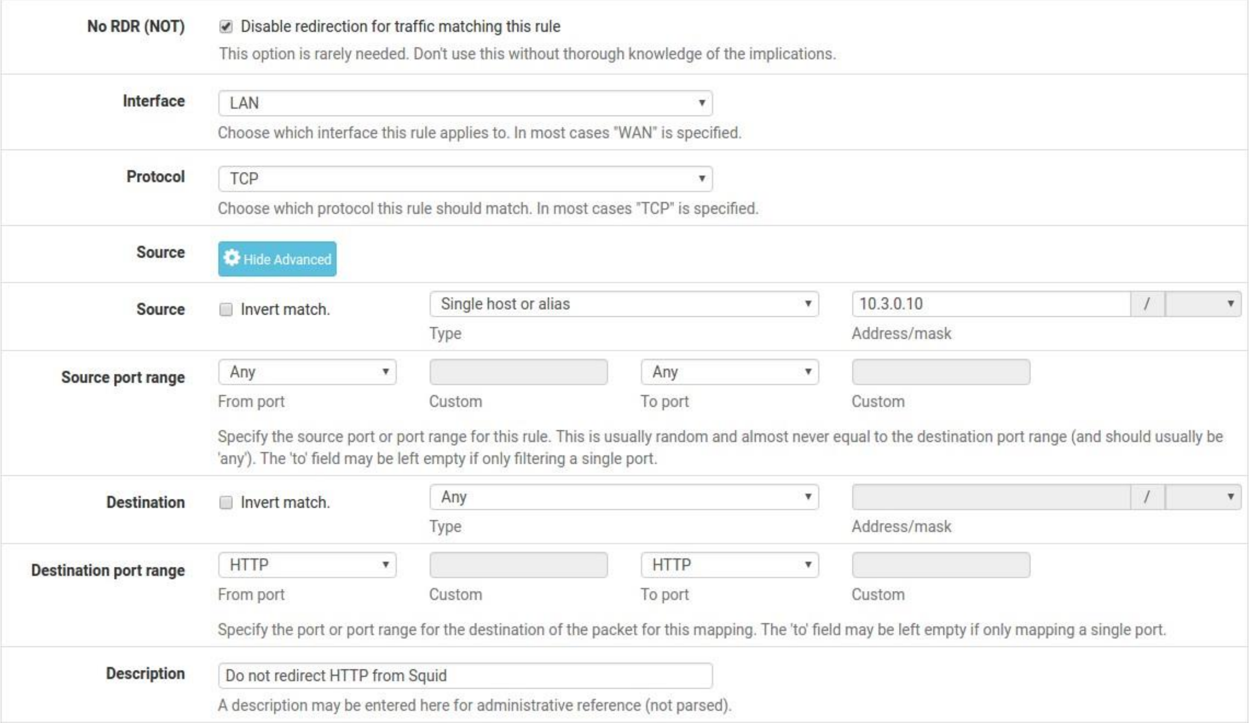

No RDR (NOT) ☑ Disable redirection for traffic matching this rule
This option is rarely needed. Don't use this without thorough knowledge of the implications.

Interface LAN
Choose which interface this rule applies to. In most cases "WAN" is specified.

Protocol TCP
Choose which protocol this rule should match. In most cases "TCP" is specified.

Source Hide Advanced

Source □ Invert match. Type: Single host or alias Address/mask: 10.3.0.10 /

Source port range From port: Any Custom: To port: Any Custom:
Specify the source port or port range for this rule. This is usually random and almost never equal to the destination port range (and should usually be 'any'). The 'to' field may be left empty if only filtering a single port.

Destination □ Invert match. Type: Any Address/mask: /

Destination port range From port: HTTP Custom: To port: HTTP Custom:
Specify the port or port range for the destination of the packet for this mapping. The 'to' field may be left empty if only mapping a single port.

Description Do not redirect HTTP from Squid
A description may be entered here for administrative reference (not parsed).

Fig. 5: Example Redirect Port Forward (Negation)

9.2 1:1 NAT

1:1 NAT (pronounced "one-to-one NAT") maps one external IPv4 address (usually public) to one internal IPv4 address (usually private). All traffic originating from that private IPv4 address going to the Internet will be mapped by 1:1 NAT to the public IPv4 address defined in the entry, overriding the Outbound NAT configuration. All traffic initiated on the Internet destined for the specified public IPv4 address on the mapping will be translated to the private IPv4 address, then evaluated against the WAN firewall ruleset. If matching traffic is permitted by the firewall rules to a target of the private IPv4 address, it will be passed to the internal host.

1:1 NAT can also translate whole subnets as well as single addresses, provided they are of the same size and align on proper subnet boundaries.

The ports on a connection remain constant with 1:1 NAT; For outbound connections, the source ports used by the local system are preserved, similar to using Static Port on outbound NAT rules.

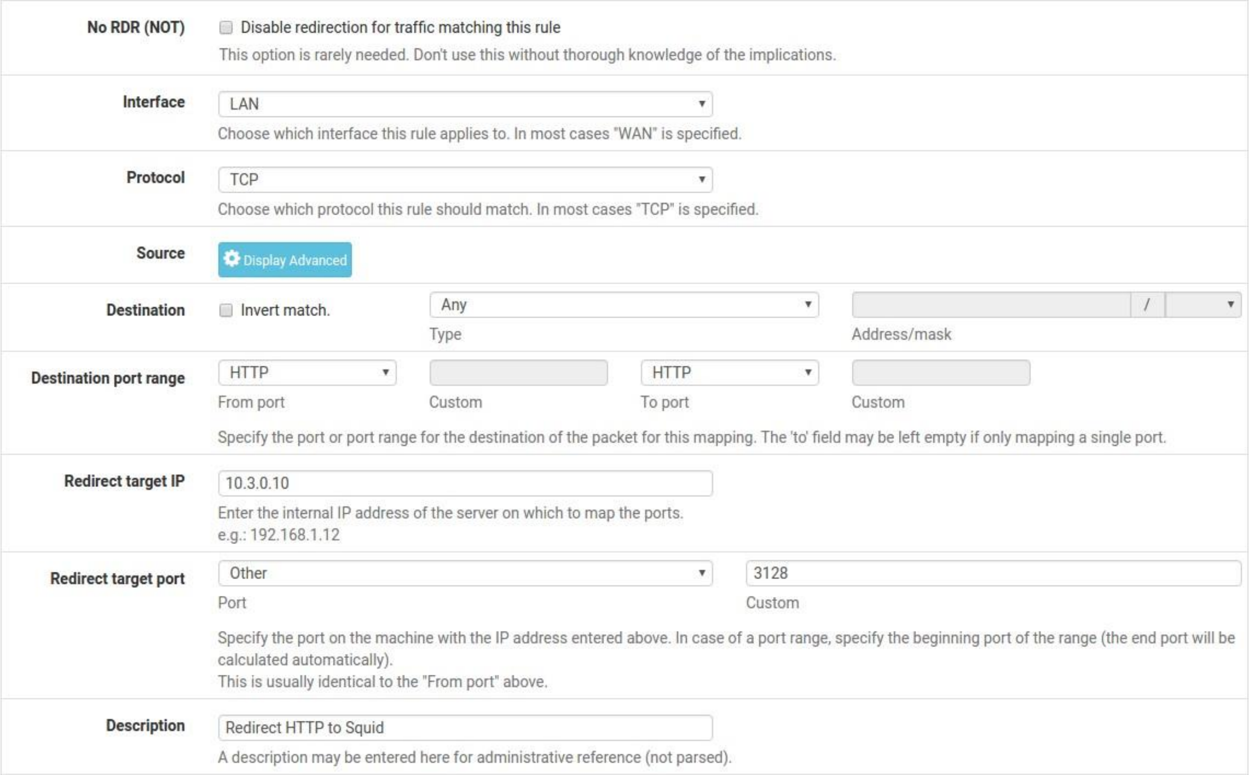

No RDR (NOT) □ Disable redirection for traffic matching this rule
This option is rarely needed. Don't use this without thorough knowledge of the implications.

Interface LAN
Choose which interface this rule applies to. In most cases "WAN" is specified.

Protocol TCP
Choose which protocol this rule should match. In most cases "TCP" is specified.

Source Display Advanced

Destination □ Invert match. Any (Type) / (Address/mask)

Destination port range HTTP (From port) (Custom) HTTP (To port) (Custom)
Specify the port or port range for the destination of the packet for this mapping. The 'to' field may be left empty if only mapping a single port.

Redirect target IP 10.3.0.10
Enter the internal IP address of the server on which to map the ports.
e.g.: 192.168.1.12

Redirect target port Other (Port) 3128 (Custom)
Specify the port on the machine with the IP address entered above. In case of a port range, specify the beginning port of the range (the end port will be calculated automatically).
This is usually identical to the "From port" above.

Description Redirect HTTP to Squid
A description may be entered here for administrative reference (not parsed).

Fig. 6: Example Redirect Port Forward

9.2.1 Risks of 1:1 NAT

The risks of 1:1 NAT are largely the same as port forwards, if WAN firewall rules permit traffic. Any time rules permit traffic, potentially harmful traffic may be admitted into the local network. There is a slight added risk when using 1:1 NAT in that firewall rule mistakes can have more dire consequences. With port forward entries, traffic is limited by constraints within the NAT rule and the firewall rule. If TCP port 80 is opened by a port forward rule, then an allow all rule on WAN would still only permit TCP 80 on that internal host. If 1:1 NAT rules are in place and an allow all rule exists on WAN, everything on that internal host will be accessible from the Internet. Misconfigurations are always a potential hazard, and this usually should not be considered a reason to avoid 1:1 NAT. Keep this fact in mind when configuring firewall rules, and as always, avoid permitting anything that is not required.

9.2.2 Configuring 1:1 NAT

To configure 1:1 NAT:

- Add a Virtual IP for the public IP address to be used for the 1:1 NAT entry as described in *Virtual IP Addresses*
- Navigate to Firewall > NAT, 1:1 tab
- Click ⤵ Add to create a new 1:1 entry at the top of the list
- Configure the 1:1 NAT entry as follows:

  Disabled Controls whether this 1:1 NAT entry is active.

  Interface The interface where the 1:1 NAT translation will take place, typically a WAN type interface.

  External subnet IP The IPv4 address to which the Internal IP address will be translated as it enters or leaves the Interface. This is typically an IPv4 Virtual IP address on Interface, or an IP address routed to the firewall via Interface.

  Internal IP The IPv4 address behind the firewall that will be translated to the External subnet IP address. This is typically an IPv4 address behind this firewall. The device with this address must use this firewall as its gateway directly (attached) or indirectly (via static route). Specifying a subnet mask here will translate the entire network matching the subnet mask. For example using x.x.x.0/24 will translate anything in that subnet to its equivalent in the external subnet.

  Destination Optional, a network restriction that limits the 1:1 NAT entry. When a value is present, the 1:1 NAT will only take effect when traffic is going from the Internal IP address to the Destination address on the way out, or from the Destination address to the External subnet IP address on the way into the firewall. The Destination field supports the use of aliases.

  Description An optional text description to explain the purpose of this entry.

  NAT reflection An override for the global NAT reflection options. *Use system default* will respect the global NAT reflection settings, *enable* will always perform NAT reflection for this entry, and *disable* will never do NAT reflection for this entry. For more information on NAT Reflection, see *NAT Reflection*.

- Click Save
- Click Apply Changes

Example Single IP Address 1:1 Configuration

This section demonstrates how to configure a 1:1 NAT entry with a single internal and external IP address. In this example, 198.51.100.210 is a Virtual IP address on the WAN interface. In most deployments this will be substituted with a working public IP addresses. The mail server in this mapping resides on a DMZ segment using internal IP address 10.3.1.15. The 1:1 NAT entry to map 198.51.100.210 to 10.3.1.15 is shown in Figure *1:1 NAT Entry*.

Example IP Address Range 1:1 Configuration

1:1 NAT can be configured for multiple public IP addresses by using CIDR ranges. In this example, 1:1 NAT is configured for a /30 CIDR range of IPs.

See also:

See *CIDR Summarization* for more information on summarizing networks or groups of IP addresses inside a larger subnet using CIDR notation.

Table 1: /30 CIDR Mapping Matching Final Octet

External IP	Internal IP
198.51.100.64/30	10.3.1.64/30
198.51.100.64	10.3.1.64
198.51.100.65	10.3.1.65
198.51.100.66	10.3.1.66
198.51.100.67	10.3.1.67

The last octet of the IP addresses need not be the same on the inside and outside, but doing so makes it logically simpler to follow. For example, Table */30 CIDR Mapping Non-Matching Final Octet* is also valid.

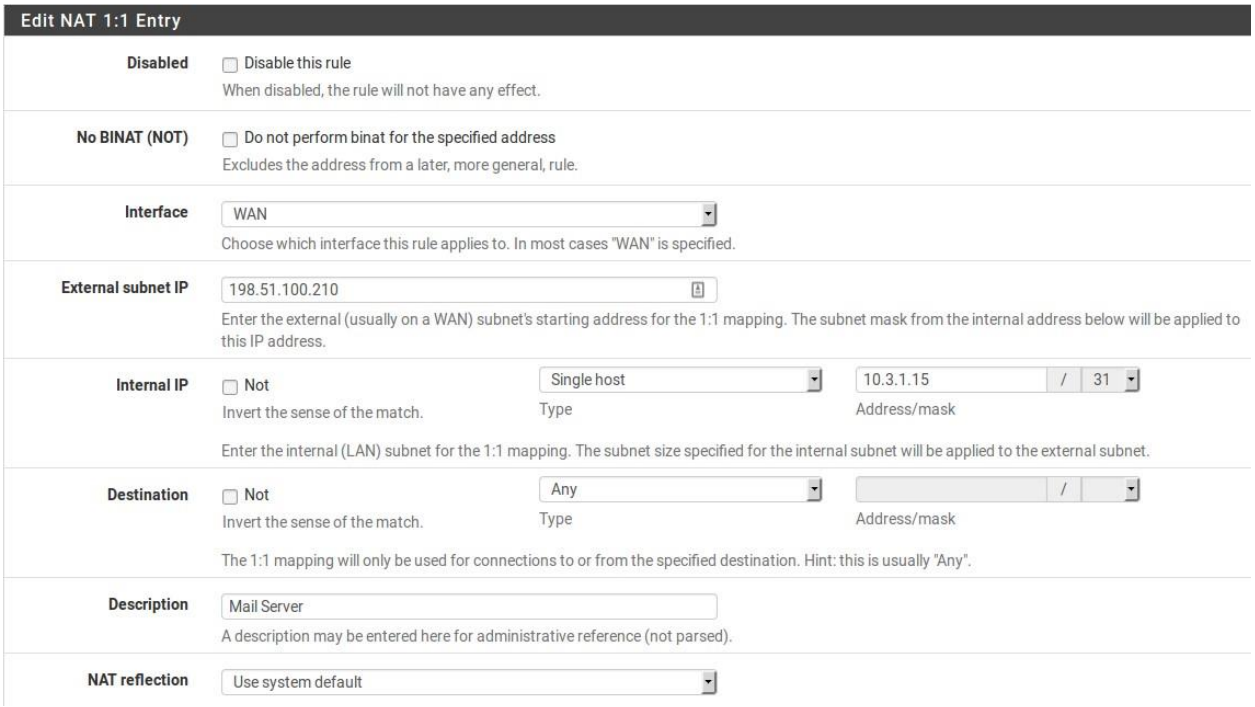

Edit NAT 1:1 Entry

Disabled □ Disable this rule
When disabled, the rule will not have any effect.

No BINAT (NOT) □ Do not perform binat for the specified address
Excludes the address from a later, more general, rule.

Interface WAN
Choose which interface this rule applies to. In most cases "WAN" is specified.

External subnet IP 198.51.100.210
Enter the external (usually on a WAN) subnet's starting address for the 1:1 mapping. The subnet mask from the internal address below will be applied to this IP address.

Internal IP □ Not — Invert the sense of the match. | Type: Single host | Address/mask: 10.3.1.15 / 31
Enter the internal (LAN) subnet for the 1:1 mapping. The subnet size specified for the internal subnet will be applied to the external subnet.

Destination □ Not — Invert the sense of the match. | Type: Any | Address/mask: /
The 1:1 mapping will only be used for connections to or from the specified destination. Hint: this is usually "Any".

Description Mail Server
A description may be entered here for administrative reference (not parsed).

NAT reflection Use system default

Fig. 7: 1:1 NAT Entry

Table 2: /30 CIDR Mapping Non-Matching Final Octet

External IP	Internal IP
198.51.100.64/30	10.3.1.200/30
198.51.100.64	10.3.1.200
198.51.100.65	10.3.1.201
198.51.100.66	10.3.1.202

198.51.100.67	10.3.1.203

Choosing an addressing scheme where the last octet matches makes the layout easier to understand and hence maintain. Figure *1:1 NAT entry for /30 CIDR range* shows how to configure 1:1 NAT to achieve the mapping listed in Table */30 CIDR Mapping Matching Final Octet*.

9.2.3 1:1 NAT on the WAN IP, aka "DMZ" on Linksys

Some consumer routers such as those from Cisco/Linksys have what they call a "DMZ" feature that will forward all ports and protocols destined to the WAN IP address to a system on the LAN. In effect, this is 1:1 NAT between the WAN IP address and the IP address of the internal system. "DMZ" in that context, however, has nothing to do with what an actual DMZ network is in real networking terminology. In fact, it's almost the opposite. A host in a true DMZ is in an isolated network away from the other LAN hosts, secured away from the Internet and LAN hosts alike. In contrast, a "DMZ" host in the Linksys meaning is not only on the same network as the LAN hosts, but completely exposed to incoming traffic with no protection.

In AZTCO-FW® software, 1:1 NAT can be active on the WAN IP address, with the caveat that it will leave all services running on the firewall itself inaccessible externally. So 1:1 NAT cannot be used on the WAN IP address in cases where VPNs of any type are enabled, or other local services on the firewall must be accessible externally. In some cases, this limitation can be mitigated by a port forward for locally hosted services.

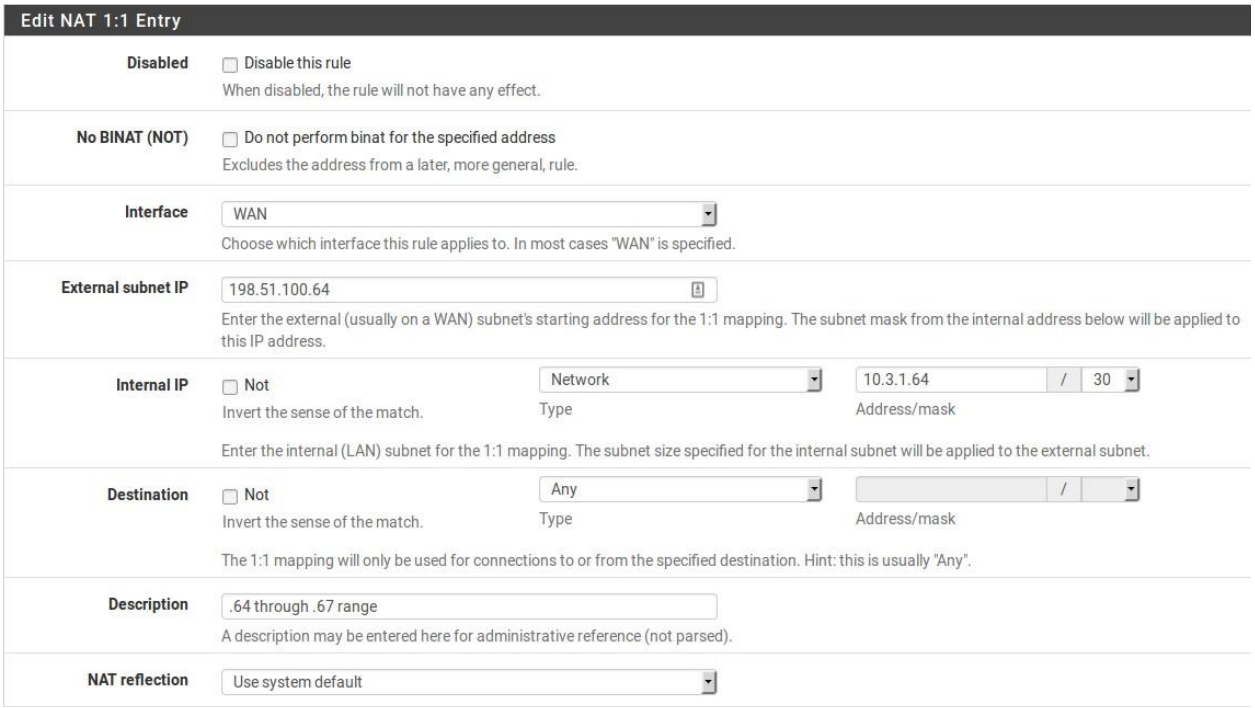

Edit NAT 1:1 Entry

Disabled: □ Disable this rule
When disabled, the rule will not have any effect.

No BINAT (NOT): □ Do not perform binat for the specified address
Excludes the address from a later, more general, rule.

Interface: WAN
Choose which interface this rule applies to. In most cases "WAN" is specified.

External subnet IP: 198.51.100.64
Enter the external (usually on a WAN) subnet's starting address for the 1:1 mapping. The subnet mask from the internal address below will be applied to this IP address.

Internal IP: □ Not — Invert the sense of the match. | Type: Network | Address/mask: 10.3.1.64 / 30
Enter the internal (LAN) subnet for the 1:1 mapping. The subnet size specified for the internal subnet will be applied to the external subnet.

Destination: □ Not — Invert the sense of the match. | Type: Any | Address/mask: /
The 1:1 mapping will only be used for connections to or from the specified destination. Hint: this is usually "Any".

Description: .64 through .67 range
A description may be entered here for administrative reference (not parsed).

NAT reflection: Use system default

Fig. 8: 1:1 NAT entry for /30 CIDR range

9.3 Ordering of NAT and Firewall Processing

Understanding the order in which firewalling and NAT occurs is important when configuring NAT and firewall rules. The basic logical order is illustrated by Figure *Ordering of NAT and Firewall Processing*. The figure also depicts where tcpdump ties in, since its use as a troubleshooting tool is described later in this documentation in *Packet Capturing*.

Each layer is not always hit in typical configurations, but the use of floating rules or manual outbound NAT or other more complicated configurations can hit each layer in both directions. The diagram only covers basic scenarios for inbound and outbound traffic.

In terms of how the ruleset is processed, the order is:

- Outbound NAT rules
- Inbound NAT rules such as Port Forwards (including rdr pass and UPnP)
- Rules dynamically received from RADIUS for IPsec and OpenVPN clients
- Internal automatic rules (pass and block for various items like lockout, snort, DHCP, etc.)
- User-defined rules:
  - Rules defined on the *floating tab*
  - Rules defined on interface group tabs (Including IPsec and OpenVPN)
  - Rules defined on interface tabs (WAN, LAN, OPTx, etc)
- Automatic VPN rules

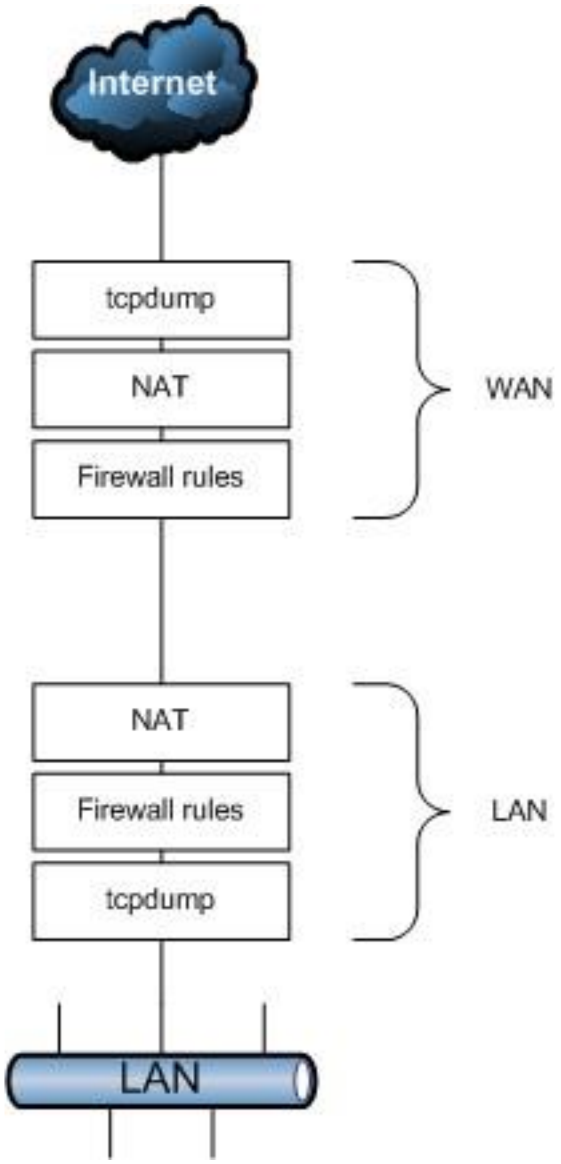

Fig. 9: Ordering of NAT and Firewall Processing

9.3.1 Firewall/NAT Processing Order Example

Traffic from LAN to WAN is processed as described in the following more detailed example. If a type of rules do not exist or do not match, they are skipped.

- Port forwards or 1:1 NAT on the LAN interface (e.g. proxy or DNS redirects)
- Firewall rules for the LAN interface:
  - Floating rules inbound on LAN
  - Rules for interface groups including the LAN interface

  - LAN tab rules

- 1:1 NAT or Outbound NAT rules on WAN
- Floating rules that match outbound on WAN

In this case, port forwards on WAN and WAN tab firewall rules do not apply.

For traffic initiated on the WAN, the order is the same but direction is reversed:

- Port forwards or 1:1 NAT on the WAN interface (e.g. public services)
- Firewall rules for the WAN interface:
  - Floating rules inbound on WAN
  - Rules for interface groups including the WAN interface
  - WAN tab rules
- 1:1 NAT or Outbound NAT rules on LAN
- Floating rules that match outbound on LAN

tcpdump is always the first and last thing to see traffic, depending on the direction. First, on the incoming interface before any NAT and firewall processing, and last on the outbound interface. It shows what is on the wire. (See *Packet Capturing*)

See also:

See *Rule Processing Order* for more information about the firewall rule processing order.

9.3.2 Floating Rules notes

Floating rules without quick set process as "last match wins" instead of "first match wins". Therefore, if a floating rule is set without quick and a packet matches that rule, then it also matches a later rule, the later rule will be used. This is the opposite of the other tab rules (groups, interfaces) and rules with quick set which stop processing as soon as a match is made. See *Floating Rules* for more details on how floating rules operate.

Extrapolating to additional interfaces

The previous diagram and lists only illustrate a basic two interface LAN and WAN deployment. When working with additional interfaces, the same rules apply. Traffic between two internal interfaces behaves the same as LAN to WAN traffic, though the default NAT rules will not translate traffic between internal interfaces so the NAT layer does not do anything in those cases. If Outbound NAT rules exist that match traffic between internal interfaces, it will apply as shown.

Rules for NAT

On the way into an interface, NAT applies before firewall rules, so if the destination is translated on the way in (e.g. port forwards or 1:1 NAT on WAN), then the firewall rules must match the translated destination. In the typical case of a port forward on WAN, this means the rule must match a destination of the target private IP address on LAN.

For example, with a port forward for TCP port 80 on WAN with an automatically added firewall rule, Figure *Firewall Rule for Port Forward to LAN Host* shows the resulting firewall rule on WAN. The internal IP address on the port forward is 10.3.0.15. Whether using port forwards or 1:1 NAT, firewall rules on all WAN interfaces must use the internal IP address as the destination.

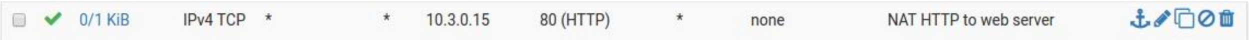

□	✓	0/1 KiB	IPv4 TCP	*	*	10.3.0.15	80 (HTTP)	*	none	NAT HTTP to web server

Fig. 10: Firewall Rule for Port Forward to LAN Host

On the way out of an interface, outbound NAT applies before firewall rules, so any floating rules matching outbound on an interface must match the source after it has been translated by outbound NAT or 1:1 NAT.

9.4 NAT Reflection

NAT reflection refers to the ability to access external services from the internal network using the external (usually public) IP address, the same as if the client were on the Internet. While many commercial and open source firewalls do not support this functionality at all, AZTCO-FW® software has good support for NAT reflection, though some environments will require a split DNS infrastructure to accommodate this functionality.

When possible, split DNS is the preferred means of accessing resources so that the firewall is not involved in accessing internal services internally. Split DNS is covered at the end of this section in *Split DNS*.

9.4.1 Configuring NAT Reflection

To enable NAT Reflection globally:

- Navigate to System > Advanced on the Firewall & NAT • Locate the Network Address

  Translation section of the page

- Configure the NAT Reflection options as follows:

  NAT Reflection mode for Port Forwards There are three available choices for NAT Reflection mode for port forwards, they are:

  Disable NAT Reflection will not be performed, but it may be enabled on a per-rule basis.

  NAT + Proxy Enables NAT Reflection using a helper program to send packets to the target of the port forward. This is useful in setups where the interface and/or gateway IP address used for communication with the target cannot be accurately determined at the time the rules are loaded. Reflection rules for use with the proxy are not created for ranges larger than 500 ports and will not be used for more than 1000 ports total between all port forwards. This mode does not work with UDP, only with TCP. Because this is a proxy, the source address of the traffic, as seen by the server, is the firewall IP address closest to the server.

Pure NAT Enables NAT Reflection using only NAT rules in pf to direct packets to the target of the port forward. It has better scalability, but it must be possible to accurately determine the interface and gateway IP address used for communication with the target at the time the rules are loaded. There are no inherent limits to the number of ports other than the limits of the protocols. All protocols available for port forwards are supported. If servers are on the same subnet as clients, the Enable automatic outbound NAT for Reflection option will mask the source of the traffic so it flows properly back through the firewall.

Reflection Timeout This option is only relevant to *NAT + Proxy* mode, and controls how long the NAT proxy daemon will wait before closing a connection. Specify the value in seconds.

Enable NAT Reflection for 1:1 NAT This option allows clients on internal networks to reach locally hosted services by connecting to the external IP address of a 1:1 NAT entry. To fully activate the feature, check both Enable NAT Reflection for 1:1 NAT and Enable automatic outbound NAT for Reflection. The latter option is only necessary if clients and servers are in the same subnet.

Enable automatic outbound NAT for Reflection When enabled, this option activates additional NAT rules for 1:1 NAT Reflection and Pure NAT mode NAT Reflection for port forwards. These additional rules mask the source address of the client to ensure reply traffic flows back through the firewall. Without this, connections between the client and server will fail as the server will reply directly back to the client using its internal IP address. The client will drop the connection since it expects a reply from the public IP address.

- Click Save to activate the new NAT reflection options

NAT Reflection Caveats

NAT reflection is a hack as it loops traffic through the firewall when it is not necessary. Because of the limited options pf allows for accommodating these scenarios, there are some limitations in the AZTCO-FW NAT + Proxy reflection implementation. Port ranges larger than 500 ports do not have NAT reflection enabled in NAT + Proxy mode, and that mode is also effectively limited to only working with TCP. The other modes require additional NAT to happen if the clients and servers are connected to the same interface of the firewall. This extra NAT hides the source address of the client, making the traffic appear to originate from the firewall instead, so that the connection can be properly established.

Split DNS is the best means of accommodating large port ranges and 1:1 NAT. Maintaining a split DNS infrastructure is required by many commercial firewalls even, and typically isn’t a problem.

9.4.2 Split DNS

A preferable alternative to NAT reflection is deploying a split DNS infrastructure. Split DNS refers to a DNS configuration where, for a given hostname, public Internet DNS resolves to public IP address, and DNS on the internal network resolves to the internal, private IP address. The means of accommodating this will vary depending on the specifics of an organization’s DNS infrastructure, but the end result is the same. NAT reflection is not necessary because hostnames resolve to the private IP addresses inside the network and clients can reach the servers directly.

Split DNS allows servers to see the true client IP address, and connections between servers and clients in the same subnet will go directly, rather than unnecessarily involving the firewall.

The only case that does not work properly with split DNS is when the external and internal port numbers are different. With split DNS, the port number has to be the same in both places.

DNS Resolver/Forwarder Overrides

If AZTCO-FW is acting as the DNS server for internal hosts, then host overrides in the DNS Resolver or DNS forwarder can provide split DNS functionality.

To add an override to the DNS Resolver:

- Navigate to Services > DNS Resolver
- Click + the under Host Overrides to reach the Host Override Options page
- Configure the host override as needed, using the internal IP address of the server. See *Host Overrides*. Figure *Add DNS Resolver Override for example.com* shows an example of a DNS override for example.com and www.example.com.
- Click Save
- Click Apply Changes

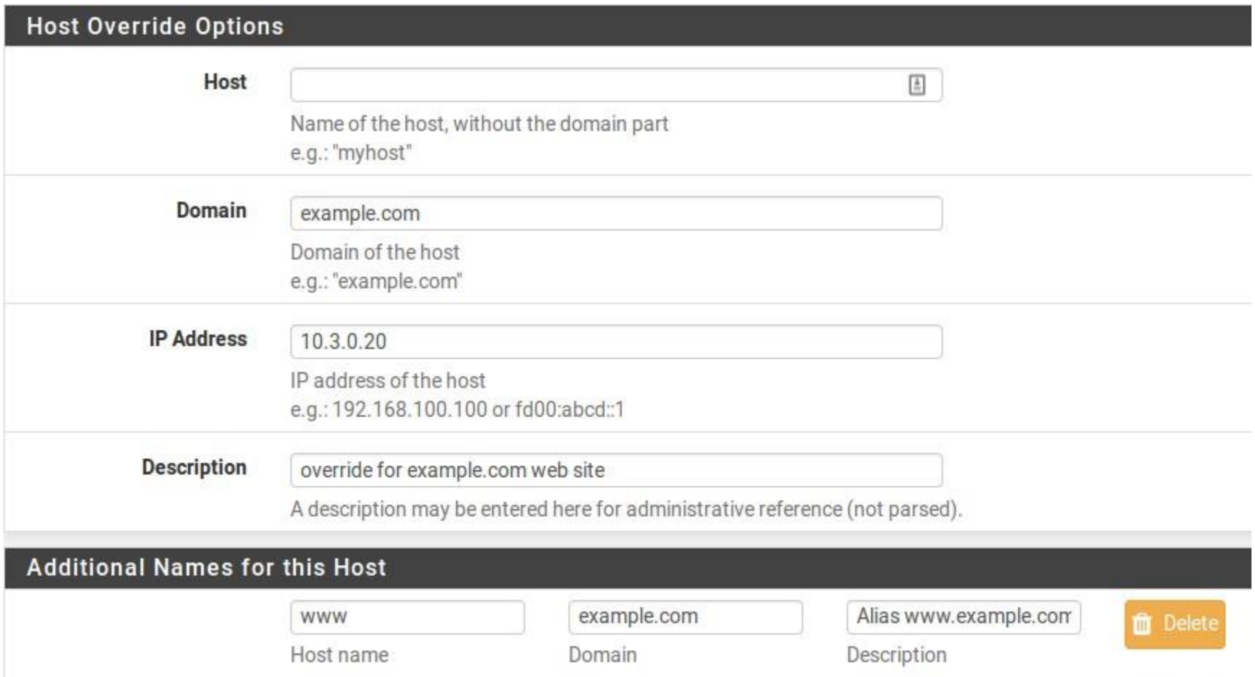

Host Override Options

Host:
Name of the host, without the domain part
e.g.: "myhost"

Domain: example.com
Domain of the host
e.g.: "example.com"

IP Address: 10.3.0.20
IP address of the host
e.g.: 192.168.100.100 or fd00:abcd::1

Description: override for example.com web site
A description may be entered here for administrative reference (not parsed).

Additional Names for this Host

Host name	Domain	Description	
www	example.com	Alias www.example.com	Delete

Fig. 11: Add DNS Resolver Override for example.com

The DNS Forwarder works identically in this regard. If the DNS Forwarder is enabled instead of the DNS Resolver, add the overrides there.

An override is required for each hostname in use behind the firewall.

Internal DNS servers

When using a separate DNS server on an internal network, such as Microsoft Active Directory, zones must be created by the DNS server administrator for all domains hosted inside the network, along with all other records for those domains (A, CNAME, MX, etc.).

In environments running the BIND DNS server where the public DNS is hosted on the same server as the private DNS, BIND's views feature is used to resolve DNS differently for internal hosts than external ones. Other DNS servers may support similar functionality. Check their documentation for information.

9.5 Outbound NAT

Outbound NAT, also known as *Source NAT*, controls how AZTCO-FW® will translate the source address and ports of traffic leaving an interface. To configure Outbound NAT, navigate to Firewall > NAT, on the Outbound tab.

There are four possible Modes for Outbound NAT:

Automatic Outbound NAT The default option, which automatically performs NAT from internal interfaces, such as LAN, to external interfaces, such as WAN.

Hybrid Outbound NAT Utilizes manual rules while also using automatic rules for traffic not matched by manually entered rules. This mode is the most flexible and easy to use for administrators who need a little extra control but do not want to manage the entire list manually.

Manual Outbound NAT Only honors the manually entered rules, and nothing more. Offers the most control, but can be tough to manage and any changes made to internal interfaces or WANs must be accounted for in the rules by hand. If the list is empty when switching from automatic to manual, the list is populated with rules equivalent to the automatically generated set.

Disable Outbound NAT Disables all outbound NAT. Useful if the firewall contains only routable addresses (e.g. public IP addresses) on all LANs and WANs.

When changing the Mode value, click the Save button to store the new value.

In networks with a single public IP address per WAN, there is usually no reason to enable manual outbound NAT. If some manual control is necessary, hybrid mode is the best choice. In environments with multiple public IP addresses and complex NAT requirements, manual outbound NAT offers more fine-grained control over all aspects of translation.

For environments using High Availability with CARP, it is important to NAT outbound traffic to a CARP VIP address, as discussed in *High Availability*. This can be accomplished in either hybrid or manual mode.

As with other rules in AZTCO-FW, outbound NAT rules are considered from the top of the list down, and the first match is used. Even if rules are present in the Outbound NAT screen, they will not be honored unless the Mode is set to Hybrid Outbound NAT or Manual Outbound NAT.

Note: Outbound NAT only controls what happens to traffic *as it leaves an interface*. It does *not* control the interface though which traffic will exit the firewall. That is handled by the routing table (*Static Routes*) or policy routing (*Policy routing*).

9.5.1 Default Outbound NAT Rules

When set to the default Automatic Outbound NAT mode, AZTCO-FW maintains a set of NAT rules to translate traffic leaving any internal network to the IP address of the WAN interface which the traffic leaves. Static route networks and remote access VPN networks are also included in the automatic NAT rules.

When outbound NAT is configured for Automatic or Hybrid modes, the automatic rules are presented in the lower section of the screen labeled Automatic Rules.

If the Outbound NAT rule list is empty, switching to Manual Outbound NAT and saving will generate a full set of rules equivalent to the automatic rules.

9.5.2 Static Port

By default, AZTCO-FW rewrites the source port on all outgoing connections except for UDP port 500 (IKE for VPN traffic). Some operating systems do a poor job of source port randomization, if they do it at all. This makes IP address spoofing easier and makes it possible to fingerprint hosts behind the firewall from their outbound traffic. Rewriting the source port eliminates these potential (but unlikely) security vulnerabilities. Outbound NAT rules, including the automatic rules, will show 🔀 in the Static Port column on rules set to randomize the source port.

Source port randomization breaks some rare applications. The default Automatic Outbound NAT ruleset disables source port randomization for UDP 500 because it will almost always be broken by rewriting the source port. Outbound

NAT rules which preserve the original source port are called Static Port rules and have ✖ on the rule in the Static Port column. All other traffic has the source port rewritten by default.

Other protocols, such as those used by game consoles, may not work properly when the source port is rewritten. To disable this functionality, use the Static Port option.

To add a rule for a device which requires static source ports:

- Navigate to Firewall > NAT, Outbound tab
- Select Hybrid Outbound NAT rule generation
- Click Save
- Click ⤵ to add a new NAT rule to the top of the list
- Configure the rule to match the traffic that requires static port, such as a source address of a PBX or a game console (See *Working with Manual Outbound NAT Rules* below)
- Check Static Port in the Translation section of the page
- Click Save
- Click Apply Changes

After making that change, the source port on outgoing traffic matching the rule will be preserved. The best practice is to use strict rules when utilizing static port to avoid any potential conflict if two local hosts use the same source port to talk to the same remote server and port using the same external IP address.

9.5.3 Disabling Outbound NAT

If public IP addresses are used on local interfaces, and thus NAT is not required to pass traffic through the firewall, disable NAT for the routable subnet. This can be achieved in several ways:

- If NAT is not required for any interface, set the outbound NAT mode to Disable
- Using Hybrid Outbound NAT, a rule set with Do not NAT can disable NAT for matching traffic
- Using Manual Outbound NAT, delete (or do not create) any NAT rules matching the routable subnets

In any of the above cases, outbound NAT will no longer be active for those source IP addresses and AZTCO-FW will then route public IP addresses without translation.

9.5.4 Working with Manual Outbound NAT Rules

Outbound NAT rules are very flexible and are capable of translating traffic in many ways.

The NAT rules are shown in a single page and the Interface column is a source of confusion for some; As traffic leaves an interface, only the outbound NAT rules set for that specific Interface are consulted.

Click ↧ from the Outbound NAT page to add a rule to the top of the list. Click ↧ to add a rule to the bottom. Place specific rules at the top, and more general rules at the bottom. The rules are processed by the firewall starting at the top of the list and working down, and the first rule to match is used. Rules may be reordered to match in the desired way.

The options for each Outbound NAT rule are:

Disabled Toggles whether or not this rule is active.

Do not NAT Checking this option causes packets matching the rule to *not* have NAT applied as they leave. This is necessary if the traffic would otherwise match a NAT rule, but must not have NAT applied. One common use for this is to add a rule exception so that the firewall IP addresses do not get NAT applied, especially in the case of CARP, where such NAT would break Internet communication from a secondary node while it is in backup mode.

Interface The interface where this NAT rule will apply when traffic is leaving via this interface. Typically this is WAN or an OPT WAN, but in some special cases it could be LAN or another internal interface.

Protocol In most cases, Outbound NAT will apply to *any* protocol, but occasionally it is necessary to restrict the protocol upon which the NAT will act. For example, to only perform static port NAT for UDP traffic from a PBX.

Source The Source is the local network which will have its address translated as it leaves the selected Interface. This is typically a LAN, DMZ, or VPN subnet. The Source Port is nearly always left blank to match all ports. This field supports the use of aliases if the Type is set to *Network*.

Note: Avoid using a source address of *any* as that will also match traffic from the firewall itself. This will cause problems with gateway monitoring and other firewall-initiated traffic.

Destination In most cases, the Destination remains set to *any* so that traffic going anywhere out of this Interface will be translated, but the Destination can be restricted as needed. For example, to translate in a certain way when going to a specific destination, such as only doing static port NAT to SIP trunk addresses. This field supports the use of aliases if the Type is set to *Network*.

Translation The Address field inside of the Translation section controls what happens to the source address of traffic matching this rule. Most commonly, this is set to *Interface Address* so the traffic

is translated to the IP address of Interface, e.g. the WAN IP address. The Address drop-down also contains all defined Virtual IP addresses, host aliases, and *Other Subnet* to manually enter a subnet for translation.

Note: An alias containing subnets cannot be used for translation. Only host aliases or a single manually entered subnet may be used.

Using a host alias or manually entered subnet, an outbound NAT rule can translate to a pool of addresses. This can help in large NAT deployments or in areas where static port is required for several clients. When translating to a host alias or subnet, a Pool Options drop-down is available with several options. Only *Round Robin* types work with host aliases. Any type may be used with a subnet.

Default Does not define any specific algorithm for selecting a translation address from the pool.

Round Robin Loops through each potential translation address in the alias or subnet in turn.

Round Robin with Sticky Address Works the same as *Round Robin* but maintains the same translation address for a given source address as long as states from the source host exist.

Random Selects a translation address for use from the subnet at random.

Random with Sticky Address Selects an address at random, but maintains the same translation address for a given source address as long as states from the source host exist.

Source Hash Uses a hash of the source address to determine the translation address, ensuring that the translated address is always the same for a given source IP address.

Bitmask Applies the subnet mask and keeps the last portion identical. For example if the source address is 10.10.10.50 and the translation subnet is 192.2.0.0/24, the rule will change the address to 192.2.0.50. This works similarly to 1:1 NAT but only in the outbound direction.

Port Specifies a specific *source* port for translation. This is almost always left blank, but could be required if the client selects a random source port but the server requires a specific source port.

Static Port Causes the original source port of the client traffic to be maintained after the source IP address has been translated. Some protocols require this, like IPsec without NAT-T, and some protocols behave better with this, such as SIP and RTP. Checking this option disables the Port entry box.

No XML-RPC Sync This option is only relevant if an HA Cluster configuration is in use, and should be skipped otherwise. When using an HA cluster with configuration synchronization, checking this box will prevent the rule from being synchronized to the other members of a cluster (see *High Availability*). Typically all rules should synchronize, however. This option is only effective on master nodes, it does *not* prevent a rule from being overwritten on slave nodes.

Description An optional text reference to explain the purpose of this rule.

These rules can accommodate most any NAT scenario, large or small.

9.5.5 Tracking Changes to Outbound NAT Rules

As mentioned in Figure *Firewall Rule Time Stamps* for firewall rules, a timestamp is added to an outbound NAT entry indicating when it was created or last edited. This timestamp shows which user created the rule, and the last person to edit the rule. When switching from Automatic Outbound NAT mode to Manual Outbound NAT mode, the created rules are marked as being created by that process.

9.6 Choosing a NAT Configuration

The best NAT configuration for a given deployment depends primarily on the number of public IP addresses available and the number of local services that require inbound access from the Internet.

9.6.1 Single Public IP Address per WAN

When only a single public IP per WAN is available, NAT options are limited. 1:1 NAT rules can be used with WAN IP addresses, but that can have drawbacks. In this case, we recommend only using port forwards.

9.6.2 Multiple Public IP Addresses per WAN

When multiple public IP addresses are available per WAN, numerous options are available for inbound and outbound NAT configuration. Port forwards, 1:1 NAT, and Hybrid or Manual Outbound NAT may all be desirable, depending on the needs of the site.

9.7 NAT and Protocol Compatibility

Some protocols do not work well with NAT and others will not work at all. Problematic protocols embed IP addresses and/or port numbers within packets (e.g. SIP and FTP), some do not work properly if the source port is rewritten (SIP from a PBX, IPsec), and some are difficult because of limitations of pf (PPTP). This section covers a sampling of protocols that have difficulties with NAT in AZTCO-FW® software, and how to work around these issues where possible.

9.7.1 FTP

FTP poses problems with both NAT and firewalls because of the design of the protocol. FTP was initially designed in the 1970s, and the current standard defining the specifications of the protocol was written in 1985. Since FTP was created more than a decade prior to NAT, and long before firewalls were common, it acts in ways that are very unfriendly toward NAT and firewalls. AZTCO-FW does not include an FTP proxy by default, but there is a client proxy available as an add-on package.

14.6. Choosing a NAT Configuration

FTP Limitations

Because pf lacks the ability to properly handle FTP traffic without a proxy, and the FTP proxy package implementation is somewhat lacking, there are some restrictions on the usage of FTP.

FTP servers behind NAT

For FTP servers behind NAT, all relevant ports must be manually forwarded in to the server and allowed in firewall rules. Or in the case of 1:1 NAT, only the firewall rules are necessary. Depending on the FTP mode, server software, and client software, some server configuration may also be required.

FTP modes

FTP can act in multiple modes that change the behavior of the client and server, and which side listens for incoming connections. The complications of NAT and firewall rules depend on these modes and whether a remote client is attempting to reach a server behind AZTCO-FW, or if a client behind AZTCO-FW is attempting to reach a remote server.

Active Mode

With Active Mode FTP, when a file transfer is requested, the client listens on a local port and then tells the server the client IP address and port. The server will then connect back to that IP address and port in order to transfer the data. This is a problem for firewalls because the port is typically random, though modern clients allow for limiting the range that is used. In the case of a client behind NAT, the IP address given would be a local address, unreachable from the server. Not only that, but a firewall rule would need to be added along with a port forward allowing traffic into this port.

When the FTP proxy package is in use and a client is behind AZTCO-FW connecting to a remote server, the proxy attempts to do three major things: First, it will rewrite the FTP PORT command so that the IP address is the WAN IP address of the firewall, and a randomly chosen port on that IP address. Next, it adds a port forward that connects the translated IP address and port to the original IP address and port specified by the FTP client. Finally, it allows traffic from the FTP server to connect to that “public” port. With Multi-WAN, the proxy will only function on the WAN containing the default gateway.

When everything is working properly, this all happens transparently. The server never knows it is talking to a client behind NAT, and the client never knows that the server isn’t connecting directly.

In the case of a server behind NAT, active mode is not usually a problem since the server will only be listening for connections on the standard FTP ports and then making outbound connections back to the clients. The outbound firewall rules must allow the server to make arbitrary outbound connections, and the rules must not policy route those connections out a WAN other than the one that accepted the inbound FTP connection.

Passive Mode

Passive Mode (PASV) acts somewhat in reverse. For clients, it is more NAT and firewall friendly because the server listens on a port when a file transfer is requested, not the client. Typically, PASV mode will work for FTP clients behind NAT without using any proxy or special handling at all.

Similar to the situation in the previous section, when a client requests PASV mode the server will provide the client with its IP address and a random port to which the client can attempt to connect. Since the server is on a private network, that IP address and port will need to be translated and allowed through the firewall. See *FTP Servers and*

Port Forwards below for rule requirements. The FTP server must provide the public IP address to which clients connect, but some clients such as Filezilla are smart enough to ignore a given IP address if it is private, and will connect to the original server IP address instead.

Extended Passive Mode

Extended Passive Mode (EPSV) works similar to PASV mode but makes allowances for use on IPv6. When a client requests a transfer, the server will reply with the port to which the client should connect. The same caveats for servers in PASV mode apply here.

FTP Servers and Port Forwards

For FTP servers providing passive mode to clients, the configuration of the FTP server must define a passive port range and must also set the external NAT address, typically the WAN IP address of the firewall. The means of setting these values varies depending on the FTP server software implementation. Consult the FTP server documentation for more information. On the firewall, the passive port range must be forwarded in with port forwards along with TCP port 21.

For FTP servers providing active mode to clients, a port forward is only required for TCP port 21.

FTP Servers and 1:1 NAT

With 1:1 NAT, firewall rules must allow port 21 and the passive port range.

9.7.2 TFTP

Standard TCP and UDP traffic initiates connections to remote hosts using a random source port in the ephemeral port range, which varies by operating system but falls within 1024-65535, and the destination port of the protocol in use. Replies from server to client reverse that: The source port is the client destination port, and the destination port is the client source port. This is how pf associates the reply traffic with connections initiated from inside a network.

TFTP (Trivial File Transfer Protocol) does not follow this convention, however. The standard defining TFTP, RFC 1350, specifies the reply from the TFTP server to client will be sourced from a pseudo-random port number. The TFTP client may choose a source port of 10325 (as an example) and use the destination port for TFTP, port 69. The server for other protocols would then send the reply using source port 69 and destination port 10325. Since TFTP instead uses a pseudo- random source port, the reply traffic will not match the state pf has created for this traffic. Hence the replies will be blocked because they appear to be unsolicited traffic from the Internet.

TFTP is not a commonly used protocol across the Internet. The only situation that occasionally comes up where this is an issue is with some IP phones that connect to outside VoIP providers on the Internet using TFTP to pull configuration and other information. Most VoIP providers do not require this.

If TFTP traffic must pass through the firewall, a TFTP proxy is available which is configured under System > Advanced on the Firewall & NAT tab. See *TFTP Proxy* for more information.

9.7.3 PPTP / GRE

The limitations with PPTP in AZTCO-FW are caused by limitations in the ability of pf to NAT the GRE protocol. As such, the limitations apply to any use of the GRE protocol, however PPTP has been the most common use of GRE in the wild.

The state tracking code in pf for the GRE protocol can only track a single session per public IP address per external server. This means if a PPTP VPN connection is in place, only one internal machine can connect simultaneously to the same a PPTP server on the Internet. A thousand machines can connect simultaneously to a thousand different PPTP servers, but only one simultaneously to a single server. A single client can also connect to an unlimited number of outside PPTP servers.

The only available work around is to use multiple public IP addresses on the firewall, one per client via Outbound or 1:1 NAT, or to use multiple public IP addresses on the external PPTP server. This is not a problem with other types of VPN connections.

Due to the extremely flawed security in PPTP (See *PPTP Warning*), including a complete compromise of the entire protocol, its usage should be discontinued as soon as possible, so this issue is not relevant given the current security standards.

9.7.4 Online Games

Games typically are NAT friendly aside from a couple caveats. This section refers to both PC games and console gaming systems with online capabilities. This section provides an overview of the experiences of numerous AZTCO-FW users. Visit the Gaming board on the AZTCO-FW forum to find more information.

Static Port

Some games do not work properly unless static port is enabled on outbound NAT rules. If a game has problems establishing a connection, the best thing to try first is enabling static port for traffic coming from the console. See *Static Port* for more information.

Multiple players or devices behind one NAT device

Some games have issues where multiple players or devices are behind a single NAT device. These issues appear to be specific to NAT, not AZTCO-FW, as users who have tried other firewalls experience the same problems with them as well. Search the Gaming board on the AZTCO-FW forum for the game or system to find information from others with similar experiences.

Overcome NAT issues with UPnP

Many modern game systems support Universal Plug-and-Play (UPnP) to automatically configure any required NAT port forwards and firewall rules. Enabling UPnP on AZTCO-FW will typically allow games to work with little or no intervention. See *UPnP & NAT-PMP* for more information on configuring and using UPnP, and for information on potential security concerns.

9.8 IPv6 Network Prefix Translation (NPt)

Network Prefix Translation, or NPt for short, works similarly to 1:1 NAT but operates on IPv6 prefixes instead. NPt can be found under Firewall > NAT on the NPt tab.

NPt takes one prefix and translates it to another. So 2001:db8:1111:2222::/64 becomes 2001:db8:3333:4444::/64 and though the prefix changes, the remainder of the address will be identical for a given host on that subnet.

Warning: NPt does NOT function like traditional outbound/overload NAT/PAT. NPt cannot be used to map an internal prefix to prefix or single address in use on a WAN, it must be used with a routed prefix.

There are a few purposes for NPt, but many question its actual usefulness. With NPt, “private” IPv6 space (fc00::/ 7) can be utilized on a LAN and it can be translated by NPt to a public, routed, IPv6 prefix as it comes and goes through a WAN. The utility of this is debatable. One use is to avoid renumbering the LAN if external providers change, however since anything external that looked for the old prefix must also be adjusted, the usefulness of that can go either way, especially when the configuration must account for avoiding collisions in the fc00::/7 space for VPN tunnels.

NPt makes perfect sense for SOHO IPv6 Multi-WAN deployments. The likelihood that a home or small business end user will have their own provider-independent IPv6 space and a BGP feed is very small. In these cases, the firewall can utilize a routed prefix from multiple WANs to function similarly to Multi-WAN on IPv4. As traffic leaves the second WAN sourced from the LAN subnet, NPt will translate it to the equivalent IP address in the routed subnet for that WAN. The LAN can either use one of the routed prefixes and do NPt on the other WANs, or use addresses in fc00::/7 and do NPt on all WANs. We recommend avoiding use of the fc00::/7 space for this task. For more information on Multi-WAN with IPv6, see *Configuring Multi-WAN for IPv6*.

When adding an NPt entry, there are few options to consider as NPt is fairly basic:

Disabled Toggles whether this rule is actively used.

Interface Selects the Interface where this NPt rule takes effect as the traffic exits.

Internal IPv6 Prefix The local (e.g. LAN) IPv6 subnet and prefix length, typically the /64 on LAN or other internal network.

Destination IPv6 Prefix The routed external IPv6 subnet and prefix length to which the Internal IPv6 Prefix will be translated. This is NOT the prefix of the WAN itself. It must be a network routed to this firewall via Interface

Description A brief description of the purpose for this entry.

Figure *NPt Example* shows an NPt rule where the LAN IPv6 subnet 2001:db8:1111:2222::/64 will be translated to 2001:db8:3333:4444::/64 as it leaves the HENETV6DSL interface.

14.8. IPv6 Network Prefix Translation (NPt)

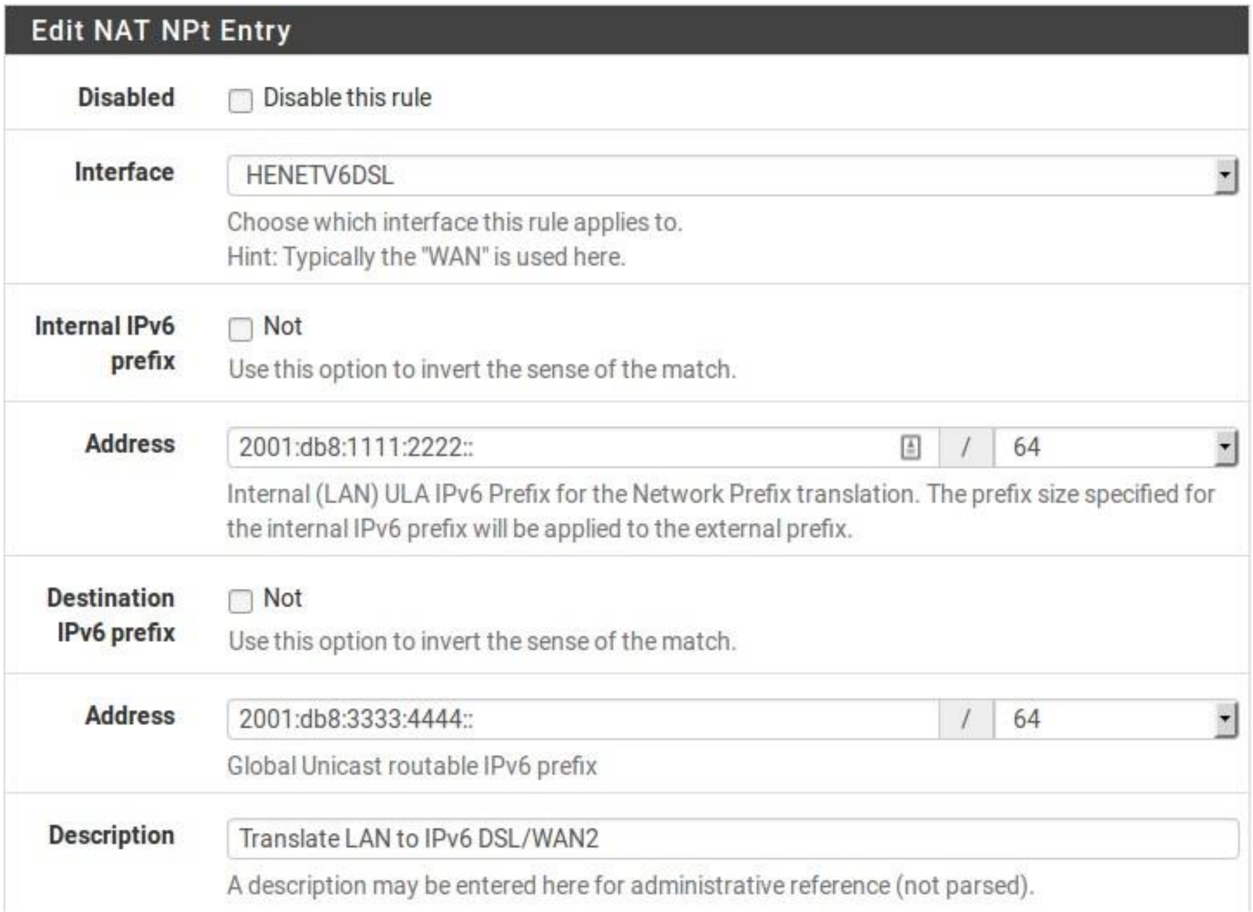

Edit NAT NPt Entry

Disabled	□ Disable this rule
Interface	HENETV6DSL Choose which interface this rule applies to. Hint: Typically the "WAN" is used here.
Internal IPv6 prefix	□ Not Use this option to invert the sense of the match.
Address	2001:db8:1111:2222:: / 64 Internal (LAN) ULA IPv6 Prefix for the Network Prefix translation. The prefix size specified for the internal IPv6 prefix will be applied to the external prefix.
Destination IPv6 prefix	□ Not Use this option to invert the sense of the match.
Address	2001:db8:3333:4444:: / 64 Global Unicast routable IPv6 prefix
Description	Translate LAN to IPv6 DSL/WAN2 A description may be entered here for administrative reference (not parsed).

Fig. 12: NPt Example

14.8. IPv6 Network Prefix Translation (NPt)

- *Troubleshooting NAT Reflection*

In its most common usage, Network Address Translation (NAT) allows multiple computers using IPv4 to be connected to the Internet using a single public IPv4 address. AZTCO-FW® software enables these simple deployments, but also accommodates much more advanced and complex NAT configurations required in networks with multiple public IP addresses.

NAT is configured in two directions: inbound and outbound. Outbound NAT defines how traffic leaving a local network destined for a remote network, such as the Internet is translated. Inbound NAT refers to traffic entering a network from a remote network. The most common type of inbound NAT is *port forwards*, which is also the type many administrators are most familiar with.

Note: In general, with the exception of Network Prefix Translation (NPt), NAT on IPv6 is not supported in AZTCO-FW. There is further discussion on the topic in *IPv6 and NAT*. Unless otherwise mentioned, this chapter is discussing NAT with IPv4.

9.9 Default NAT Configuration

This section describes the default NAT configuration present on AZTCO-FW. The most appropriate NAT configuration that can be determined is generated automatically. In some environments, this configuration may not be suitable, and AZTCO-FW fully enables changing it from the web interface. This is a contrast from many other open source firewall distributions, which do not allow the capabilities commonly required in all but small, simple networks.

9.9.1 Default Outbound NAT Configuration

In a typical two-interface AZTCO-FW setup with LAN and WAN, the default NAT configuration automatically translates Internet-bound traffic to the WAN IP address. When multiple WAN interfaces are configured, traffic leaving any WAN interface is automatically translated to the address of the WAN interface being used.

Static port is automatically configured for IKE (part of IPsec). Static port is covered in more detail in *Outbound NAT* about Outbound NAT.

For detecting WAN-type interfaces for use with NAT, AZTCO-FW looks for the presence of a gateway selected on the interface configuration if it has a static IP address, or AZTCO-FW assumes the interface is a WAN if it is a dynamic type such as PPPoE or DHCP.

9.9.2 Default Inbound NAT Configuration

By default, nothing is allowed in from the Internet on the WAN interface. If traffic initiated on the Internet must be allowed to reach a host on the internal network, port forwards or 1:1 NAT are required. This is covered in the coming sections.

CHAPTER

TEN

ROUTING

10.1 Gateways

Gateways are the key to routing; They are routers through which other networks can be reached. The kind of gateway most people are familiar with is a *default* gateway, which is the router through which a host will communicate to the Internet or any other networks it doesn't have a more specific route to reach. Gateways are also used for static routing, where other networks must be reached via specific local routers. On most networks, gateways reside in the same subnet as one of the interfaces on a host. For example, if a firewall has an IP address of 192.168.22.5/24, then a gateway to another network would have to be somewhere inside of 192.168.22.x if the other network is reachable through that interface. One notable exception to this is point-to-point interfaces like those used in PPP-based protocols, which often have gateway IP addresses in another subnet because they are not used in the same way.

10.1.1 Gateway Address Families (IPv4 and IPv6)

When working with routing and gateways, the functionality and procedures are the same for both IPv4 and IPv6 addresses, however all of the addresses for a given route must involve addresses of the same family. For example, an IPv6 network must be routed using an IPv6 gateway/router. A route cannot be created for an IPv6 network using an IPv4 gateway address. When working with gateway groups, the same restriction applies; All gateways in a gateway group must be of the same address family.

10.1.2 Managing Gateways

Before a gateway can be utilized for any purpose, it must be added to the firewall configuration.

If a gateway will be used for a WAN-type interface, it can be added on the configuration page for that interface (See *Interface Configuration Basics*), or it may be added first manually and then selected from the drop-down list on the interface configuration.

Dynamic interface types such as DHCP and PPPoE receive an automatic gateway that is noted as Dynamic in the gateway list. The parameters for such gateways can be adjusted the same as the parameters for a static gateway.

Note: Deleting a dynamic gateway will clear its custom settings, but the dynamic gateway itself cannot be removed.

To add or manage gateways, navigate to System > Routing, Gateways tab.

On the screen there are a variety of options to manage gateway entries:

- ✚ Add at the bottom of the list creates a new gateway
- ⊘ creates a copy of an existing gateway
- ⧉ edits an existing gateway
- 🗑 deletes a gateway
- ☑ disables an active gateway
- 🗑 enables a disabled gateway

The individual options for gateways are discussed in detail in *Gateway Settings*.

Managing the Default Gateway

The Default Gateway section at the bottom of System > Routing, Gateways tab controls which gateway(s) are used by default when the firewall routes traffic. Traffic from the firewall itself will follow the default gateway, as will traffic passing through the firewall when it does not match other more specific routes or policy routing rules.

There are two controls in the section which set the default gateway for IPv4 and IPv6 respectively.

The default gateway can have one of the following values:

Automatic The firewall will automatically use gateways from this list (from the top down) for the default gateway, switching to the next item in the list if gateways fail or are marked down.

For more control over this behavior, use a gateway group instead.

Gateway The selected single gateway is always used for the default gateway.

Gateway Group The firewall uses the selected *gateway group* pick the default gateway. It will change from one gateway to another if the preferred default fails.

> Warning: This function does not support load balancing, only failover. When using a gateway group for the default gateway, the group must only have one gateway in each tier.

None No default gateway for the address family will be added to the routing table.

10.2 Gateway Settings

When adding or editing a gateway, the GUI presents a page with the options for controlling gateway behavior.

The only required settings are the Interface, Address Family, Name, and the Gateway (IP address).

Interface The interface through which the gateway is reached. For example, if this is a local gateway on the LAN subnet, choose the *LAN* interface here.

Address Family Either *IPv4* or *IPv6*, depending on the type of address for this gateway.

Name The Name for the gateway, as referenced in the gateway list, and various drop-down and other selectors for gateways. It can only contain alphanumeric characters, or an underscore, but no spaces. For example: WANGW, GW_WAN, and WANGATE are valid but WAN GW is not allowed.

Gateway The IP address of the gateway. As mentioned previously, this must reside in a subnet directly configured on the selected Interface.

Note: Rare cases which require a gateway not in the interface subnet can still function, but require an additional setting. See the Use Non-Local Gateway option in *Advanced Gateway Settings* for details.

Disable Gateway Monitoring By default, the gateway monitoring daemon will ping each gateway periodically to monitor latency and packet loss for traffic to the monitored IP address. This data is used for gateway status information and also to draw the Quality RRD graph. If this monitoring is undesirable for any reason, it may be disabled by checking Disable Gateway Monitoring. Note that if the gateway status is not monitored, then Multi-WAN will not work properly as it cannot detect failures.

Disable Gateway Monitoring Action When set, the gateway monitoring daemon will take no action if the status of the gateway changes. For example, no events will be acted upon if it becomes unresponsive or suffers from high latency.

This is useful if the administrator wants to monitor a gateway without the monitoring causing additional disruptions.

Monitor IP The Monitor IP address option configures the IP address used by the gateway monitoring daemon to determine the gateway status using ICMP echo requests (“pings”).

By default the gateway monitoring daemon will ping the gateway IP address. This is not always desirable, especially in the case where the gateway IP address is local, such as on a cable modem or fiber CPE. In those cases it makes more sense to ping something farther upstream, such as an ISP DNS server or a server on the Internet. Another case is when an ISP is prone to upstream failures, so pinging a host on the Internet is a more accurate test to determine if a WAN is usable rather than testing the link itself. Some popular choices include Google public DNS servers, or popular web sites such as Google or Yahoo. If the IP address specified in this box is not directly connected, a static route is added to ensure that traffic to the Monitor IP address leaves via the expected gateway. Each gateway must have a unique Monitor IP address.

The status of a gateway as perceived by the firewall can be checked by visiting Status > Gateways or by using the Gateways widget on the dashboard. If the gateway shows Online, then the monitor IP address is successfully responding to pings.

Force State When Mark Gateway as Down is checked, the gateway will always be considered down, even when pings are returned from the monitor IP address. This is useful for cases when a WAN is behaving inconsistently and the gateway transitions are causing disruption. The gateway can be forced into a *down* state so that other gateways may be preferred until it stabilizes.

Description An optional Description of the gateway entry for reference. A short note about the gateway or interface it's used for may be helpful, or it may be left blank.

10.2.1 Advanced Gateway Settings

Several parameters can be changed to control how a gateway is monitored or treated in a Multi-WAN scenario. Most users will not need to alter these values. To access the advanced options, click the ⚙ Display Advanced button. If any of the advanced options are set, this section is automatically expanded. For more information on using multiple WAN connections, see *Multiple WAN Connections*.

Weight When using Multi-WAN load balancing, if two gateways have different amounts of bandwidth, the Weight parameter adjusts the ratio at which connections are sent through each gateway.

For example if WAN1 has 50Mbit/s and WAN2 has 100Mbit/s, weight WAN1 as *1* and WAN2 as *2*. Then for every three connections that go out, one will use WAN1 and two will use WAN2. Using this method, connections are distributed in a way that is more likely to better utilize the available bandwidth. Weight from *1* to *30* may be chosen.

Data Payload To conserve bandwidth, the dpinger daemon sends a ping with a payload size of 0 by default so that no data is contained within the ICMP echo request. However, in rare circumstances a CPE, ISP router, or intermediate hop may drop or reject ICMP packets without a payload. In these cases, set the payload size above 0. Usually a size of 1 is enough to satisfy affected equipment.

Latency Thresholds The Latency Thresholds fields control the amount of latency that is considered normal for this gateway. This value is expressed in milliseconds (ms). The default values are From 200 and To 500.

The value in the From field is the lower boundary at which the gateway would be considered in a warning state, but not down. If the latency exceeds the value in the To field, it is considered down and removed from service. The proper values in these fields can vary depending on what type of connection is in use, and what ISP or equipment is between the firewall and the monitor IP address.

Some common situations may require adjusting these values. For instance some DSL lines operate acceptably even at higher latency, so increasing the To parameter to 700 or more would lower the number of times the gateway would be considered down when, in fact, it was passing traffic sufficiently. Another example is a GIF tunnel to a provider such as he.net for IPv6. Due to the nature of GIF tunnels and load on the tunnel servers, the tunnel could be working acceptably even with latency as high as 900 ms as reported by ICMP ping responses.

Packet Loss Thresholds Similar to Latency Thresholds, the Packet Loss Thresholds control the amount of packet loss to a monitor IP address before it would be considered unusable. This value is expressed as a percentage, 0 being no loss and 100 being total loss. The default values are From 10 and To 20.

The value in the From field is the lower boundary at which the gateway would be considered in a warning state, but not down. If the amount of packet loss exceeds the value in the To field, it is considered down and removed from service. The proper values in these fields can vary depending on what type of connection is in use, and what ISP or equipment is between the firewall and the monitor IP address.

As with latency, connections can be prone to different amounts of packet loss and still function in a usable way, especially if the path to a monitor IP address drops or delays ICMP in favor of other traffic. We have observed unusable connections with minor amounts of loss, and some that are usable even when showing 45% loss. If loss alarms occur on a normally functioning WAN gateway, enter higher values in the From and To fields until a good balance for the circuit is achieved.

Probe Interval The value in the Probe Interval field controls how often a ping is sent to the monitor IP address, in *milliseconds*. The default is to ping twice per second (500 ms).

In some situations, such as links that need monitored but have high data charges, even a small ping every second can add up. This value can be safely increased so long as it less than or equal to the

Alert Interval and also does not violate the constraint on the Time Period listed below. Lower values will ping more often and be more accurate, but consume more resources. Higher values will be less sensitive to erratic behavior and consume less resources, at the cost of accuracy.

Note: The quality graph is averaged over seconds, not intervals, so as the Probe Interval is increased the accuracy of the quality graph is decreased.

Loss Interval Time in milliseconds before packets are treated as lost. The default is 2000 ms (2 seconds). Must be greater than or equal to the High Latency Threshold.

If a circuit is known to have high latency while operating normally, this can be increased to compensate.

Time Period The amount of time, in milliseconds, over which ping results are averaged. The default is 60000 (60 seconds, one minute). A longer Time Period will take more time for latency or loss to trigger an alarm, but it is less prone to be affected by erratic behavior in ping results.

The Time Period must be greater than twice the sum of the Probe Interval and Loss Interval, otherwise there may not be at least one completed probe.

Alert Interval The time interval, in milliseconds, at which the daemon checks for an alert condition. The default value is 1000 (1 second). This value must be greater than or equal to the Probe Interval, because an alert could not possibly occur between probes.

Use Non-Local Gateway The Use non-local gateway through interface specific route option allows a non-standard configuration where a gateway IP address exists outside of an interface subnet. Some providers attempting to scrape the bottom of the IPv4 barrel have resorted to this in order to not put a gateway into each customer subnet. Do not activate this option unless required to do so by the upstream provider.

10.3 Gateway Groups

Gateway groups are a set of gateways, but are treated as one entity in gateway fields of the GUI. Groups will appear in the gateway drop-downs available on, for example, firewall rule editing.

Gateway groups are managed from the Groups tab on System > Routing.

10.3.1 Gateway Group Options

When creating a gateway group, the following options are available:

Group Name The name of this gateway group. The name must be less than 32 characters in length, and may only contain letters a-z, digits 0-9, and an underscore. This will be the name used to refer to this gateway group in the Gateway field in firewall rules. This field is required.

Gateway Priority This list contains every gateway on the firewall to select which gateways will be a part of this group. The GUI will filter the list address family after the first selection.

Tier The priority level for this gateway. The value may be from 1-5 or *Never* to exclude the gateway from this group.

Lower values are higher priority. For example, gateways on *Tier 1* are used before gateways on *Tier 2*, and so on.

15.3. Gateway Groups

Gateways on the same tier are used by the firewall for load balancing when possible. Load balancing naturally performs failover as failed gateways are removed from the pool available for load balancing.

Gateways on different tiers result in failover from gateways on lower tiers to those higher tiers. For example, if Tier 1 contains only one gateway and it fails, then the next tier (Tier 2) is checked for available gateways and the firewall uses those instead, and so on.

> Warning: Some firewall features which support gateway groups only support failover, not load balancing. For example, when using a gateway group for the default gateway or as a VPN endpoint, each gateway must be on a separate tier.

Virtual IP When using a gateway group for failover in certain contexts which require binding a specific address, such as IPsec, this option controls which address on an interface is used for that purpose. For example, in an HA pair this could be a CARP VIP used as an endpoint for IPsec tunnels.

Leave it set to the default *Interface Address* when a specific address is not required by any use of the gateway group.

Trigger Level

Configures how the firewall manages the gateway group entries when certain types of gateway events occur.

Member Down Marks the gateway as down only when it is completely down, past one or both of the higher thresholds configured for the gateway. This catches the worst sort of failures, when the gateway is completely unresponsive, but may miss more subtle issues with the circuit that can make it unusable long before the gateway reaches that level.

Packet Loss Marks the gateway as down when packet loss crosses the lower alert threshold (See *Advanced Gateway Settings*).

High Latency Marks the gateway as down when latency crosses the lower alert threshold (See *Advanced Gateway Settings*).

Packet Loss or High Latency Marks the gateway as down for either type of alert.

Description Text describing the purpose of this gateway group.

10.3.2 Tier Priority Example

Example:

- WANGW: Tier1
- OPT1GW: Tier2
- OPT3GW: Tier3

In the example above OPT1GW would be used if WANGW fails, OPT3GW will be used if both WANGW and OPT1GW fail.

10.3.3 Connection-Based Round-Robin Load Balancing Example

Example:

- WANGW: Tier1
- OPT1GW: Tier1
- OPT3GW: Tier1

In the example above all gateways have the same Tier value. When this group is used by a firewall rule, connections matching that rule will perform connection-based round-robin load balancing between all of the gateways.

Note: If any of the gateways fail, they are automatically removed from active usage in the group, effectively resulting in failover in addition to load balancing.

10.4 Static Routes

Static routes are used when hosts or networks are reachable through a router other than the default gateway. AZTCO-FW® knows about the networks directly attached to it, and reaches all other networks as directed by its routing table. In networks where an internal router connects additional internal subnets, a static route must be defined for that network to be reachable. The routers through which these other networks are reached must first be added as gateways. See *Gateways* for information on adding gateways.

Static routes are found under System > Routing on the Routes tab.

See also:

- *Accessing Firewall Services over IPsec VPNs*
- *Policy Routing Configuration*

10.4.1 Static Route Configuration

When adding or editing a static route, the following options are available:

Destination Network The network and subnet mask reachable using this route. This may be an IPv4 address (subnet ID), IPv6 prefix, or an *alias*.

Gateway The router through which this network is reached.

Disabled Check if the static route should not be used, only defined.

Description Text to describe the route, its purpose, etc.

10.4.2 Managing Static Routes

To add a route:

- Navigate to System > Routing on the Routes tab
- Click + Add to create a new static route
- Fill in the configuration as described in *Static Route Configuration*
- Click Save
- Click Apply Changes

To manage existing routes, navigate to System > Routing on the Routes tab. On the screen there are a variety of options to manage routes:

- ⧉ edits an existing route
- ⊘ creates a copy of an existing gateway
- 🗑 •deletes a route
- ☑ •disables an active route
- 🗑 •enables a disabled route

10.4.3 Example Static Route

Figure *Static Routes* illustrates a scenario where a static route is required.

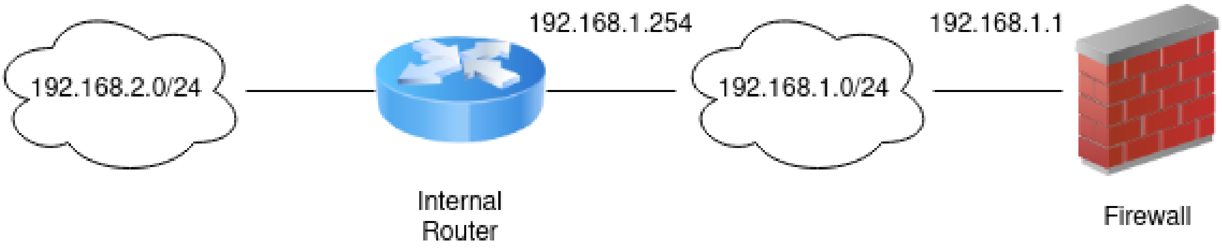

Fig. 1: Static Routes

Because the 192.168.2.0/24 network in Figure *Static Routes* is not on an interface directly connected to the firewall, a static route is required for the firewall to know how to reach that network. Figure *Static Route Configuration* shows the appropriate static route for the above diagram. As mentioned earlier, before a static route may be added a gateway must first be defined.

LAN firewall rules must allow traffic to pass from a source of the networks reachable via static routes on LAN, and outbound NAT must also accommodate these networks.

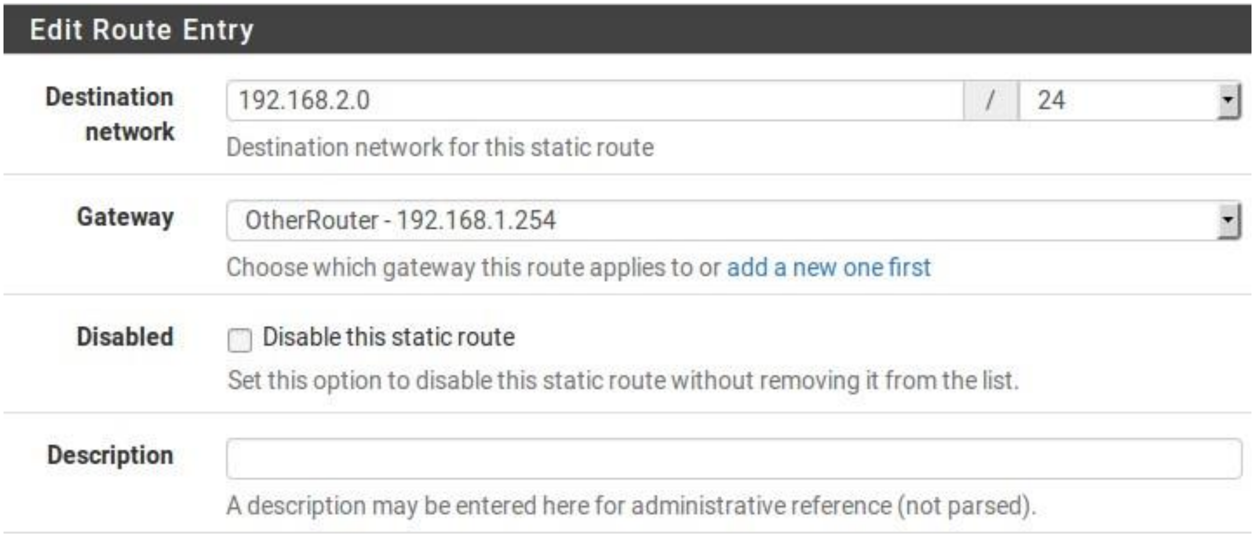
Edit Route Entry

Destination network: 192.168.2.0 / 24
Destination network for this static route

Gateway: OtherRouter - 192.168.1.254
Choose which gateway this route applies to or add a new one first

Disabled: □ Disable this static route
Set this option to disable this static route without removing it from the list.

Description:
A description may be entered here for administrative reference (not parsed).

Fig. 2: Static Route Configuration

10.4.4 Bypass Firewall Rules for Traffic on Same Interface

In some environments using static routes, traffic ends up routing asymmetrically. This means the traffic will follow a different path in one direction than the traffic flowing in the opposite direction. Take Figure *Asymmetric Routing* for example.

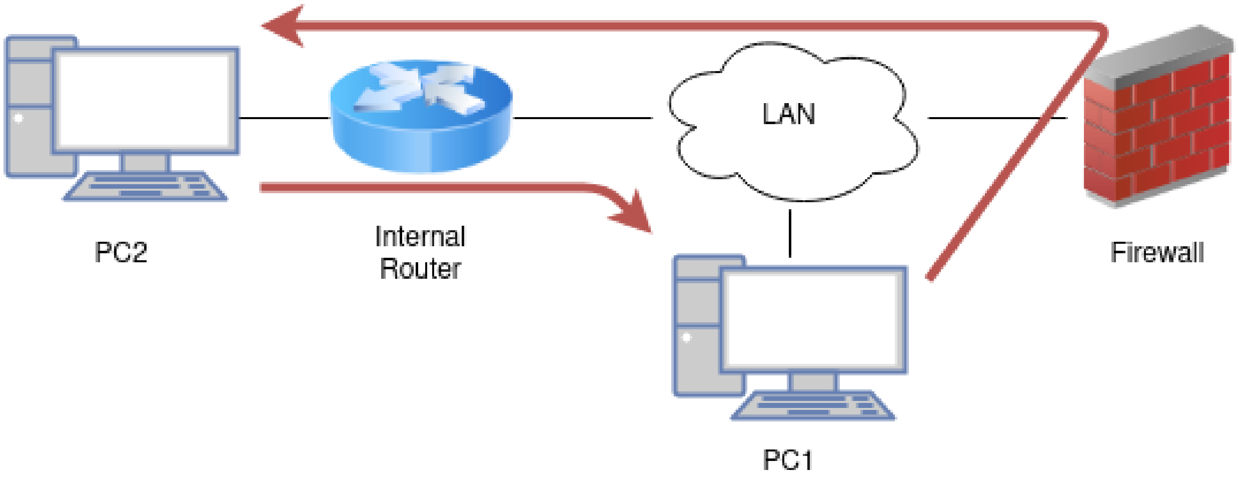

Fig. 3: Asymmetric Routing

Traffic from PC1 to PC2 will go through the firewall since it is the default gateway for PC1, but traffic in the opposite direction will go directly from the router to PC1.

Since AZTCO-FW is a stateful firewall, it must see traffic for the entire connection to be able to filter traffic properly. With asymmetric routing such as in this example, any stateful firewall will drop legitimate traffic because it cannot properly keep state without seeing traffic in both directions. This generally only affects TCP, since other protocols do not have a formal connection handshake the firewall can recognize for use in state tracking.

Note: A connection may appear to work for a short time and then stop. This can be from the firewall removing a state which doesn't transition to a fully open state, or it may be from clients expiring an ICMP redirect.

In asymmetric routing scenarios, there is an option in the firewall GUI which can be used to prevent legitimate traffic from being dropped. The option adds firewall rules which allow all traffic between networks defined in static routes using a more permissive set of rule options and state handling. To activate this option:

- Click System > Advanced
- Click the Firewall/NAT tab
- Check Bypass firewall rules for traffic on the same interface
- Click Save

Alternatively, firewall rules may be added manually to allow similar traffic. Two rules are needed, one on the interface tab where the traffic enters (e.g. LAN) and another on the Floating tab:

- Navigate to Firewall > Rules
- Click the tab for the interface where the traffic will enter (e.g. LAN)
- Click ⤵ Add to add a new rule to the top of the list
- Use the following settings:

  Protocol *TCP*

Source The local systems utilizing the static route (e.g. *LAN Net*)

Destination The network on the other end of the route

TCP Flags Check Any flags (Under Advanced Features)

State Type *Sloppy State* (Under Advanced Features)

- Click Save
- Click the Floating tab
- Click ↓ Add to add a new rule to the top of the list
- Use the following settings:

Interface The interface where the traffic originated (e.g. LAN)

Direction *Out*

Protocol *TCP*

Source The local systems utilizing the static route (e.g. *LAN Net*)

Destination The network on the other end of the route

TCP Flags Check Any flags (Under Advanced Features)

State Type *Sloppy State* (Under Advanced Features)

- Click Save

If additional traffic from other sources or destinations is shown as blocked in the firewall logs with TCP flags such as "TCP:SA" or "TCP:PA", the rules may be adjusted or copied to match that traffic as well.

Note: If filtering of traffic between statically routed subnets is required, it must be done on the router and not the firewall since the firewall is not in a position on the network where it can effectively control that traffic.

See also:

- *Troubleshooting Asymmetric Routing*

10.4.5 ICMP Redirects

When a device sends a packet to its default gateway, and the gateway knows the sender can reach the destination network via a more direct route, it will send an ICMP redirect message in response and forward the packet as configured. The ICMP redirect causes a route for that destination to be temporarily added to the routing table of the sending device, and the device will subsequently use that more direct route to reach that destination.

This will only work if the client OS is configured to permit ICMP redirects, which is typically the case by default.

ICMP redirects are common when static routes are present which point to a router on the same interface as client PCs and other network devices. The asymmetric routing diagram from the previous section is an example of this.

ICMP redirects have a mostly undeserved bad reputation from some in the security community because they allow modification of a client routing table. However they are not the risk that some imply, as to be accepted, the ICMP redirect message must include the first 8 bytes of data from the original datagram. A host in a position to see that data and hence be able to successfully forge illicit ICMP redirects is in a position to accomplish the same end result in multiple other ways.

See also:

- *Route Table Contents*
- *Multiple WAN Connections*
- *IPv6 Router Advertisements*
- *OpenBGPD package*
- *Routing Information Protocol (RIP)*
- *Routing Public IP Addresses*
- *Dynamic Routing Protocol Basics*
- *Troubleshooting Gateway Monitoring*
- *Troubleshooting "No buffer space available" Errors*
- *Troubleshooting Network Connectivity*
- *Troubleshooting Traceroute Output*
- *Troubleshooting Website Access*
- *Troubleshooting Routes*

One of the primary functions of a firewall is routing traffic. This chapter covers several topics related to routing, including gateways, static routes, routing protocols, routing of public IP addresses, and displaying routing information.

CHAPTER

ELEVEN

BRIDGING

11.1 Creating a Bridge

In AZTCO-FW® software, bridges are added and removed at Interfaces > Assignments on the Bridges tab. Using bridges, any number of ports may be bound together easily. Each bridge created in the GUI will also create a new bridge interface in the operating system, named bridgeX where X starts at 0 and increases by one for each new bridge. These interfaces may be assigned and used like most other interfaces, which is discussed later in this chapter.

To create a bridge:

- Navigate to Interfaces > Assignments on the Bridges tab.
- Click Add to create a new bridge.
- Select at least one entry from Member Interfaces. Select as many as needed using Ctrl-click.
- Add a Description if desired.
- Click Show Advanced Options to review the remaining configuration parameters as needed. For most cases they are unnecessary.
- Click Save to complete the bridge.

Note: A bridge may consist of a single member interface, which can help with migrating to a configuration with an assigned bridge, or for making a simple span/mirror port.

11.2 Advanced Bridge Options

There are numerous advanced options for a bridge and its members. Some of these settings are quite involved, so they are discussed individually in this section.

11.2.1 (Rapid) Spanning Tree Options

Spanning Tree is a protocol that helps switches and devices determine if there is a loop and cut it off as needed to prevent the loop from harming the network. There are quite a few options that control how spanning tree behaves which allow for certain assumptions to be made about specific ports or to ensure that certain bridges get priority in the case of a loop or redundant links. More information about STP may be found in the FreeBSD ifconfig(8) man page, and on Wikipedia.

Protocol

The Protocol setting controls whether the bridge will use IEEE 802.1D Spanning Tree Protocol (*STP*) or IEEE 802.1w Rapid Spanning Tree Protocol (*RSTP*). RSTP is a newer protocol, and as the name suggests it operates much faster than STP, but is backward compatible. The newer IEEE 802.1D-2004 standard is based on RSTP and makes STP obsolete.

Select STP only when older switch gear is in use that does not behave well with RSTP.

STP Interfaces

The STP Interfaces list reflects the bridge members upon which STP is enabled. Ctrl-click to select bridge members for use with STP.

Valid Time

Set the Valid Time for a Spanning Tree Protocol configuration. The default is 20 seconds. The minimum is 6 seconds and the maximum is 40 seconds.

Forward Time

The Forward Time option sets the time that must pass before an interface begins forwarding packets when Spanning Tree is enabled. The default is 15 seconds. The minimum is 4 seconds and the maximum is 30 seconds.

Note: A longer delay will be noticed by directly connected clients as they will not be able to pass traffic, even to obtain an IP address via DHCP, until their interface enters forwarding mode.

Hello Time

The Hello Time option sets the time between broadcasting of Spanning Tree Protocol configuration messages. The

Hello Time may only be changed when operating in legacy STP mode. The default is 2 seconds. The minimum is 1

second and the maximum is 2 seconds. **16.2. Advanced Bridge Options**

Bridge Priority

The Bridge Priority for Spanning Tree controls whether or not this bridge would be selected first for blocking should a loop be detected. The default is 32768. The minimum is 0 and the maximum is 61440. Values must be a multiple of 4096. Lower priorities are given precedence, and values lower than 32768 indicate eligibility for becoming a root bridge.

Hold Count

The transmit Hold Count for Spanning Tree is the number of packets transmitted before being rate limited. The default is 6. The minimum is 1 and the maximum is 10.

Port Priorities

The Priority fields set the Spanning Tree priority for each bridge member interface. Lower priorities are given preference when deciding which ports to block and which remain forwarding. Default priority is 128, and must be between 0 and 240.

Path Costs

The Path Cost fields sets the Spanning Tree path cost for each bridge member. The default is calculated from the link speed. To change a previously selected path cost back to automatic, set the cost to 0. The minimum is 1 and the maximum is 200000000. Lower cost paths are preferred when making a decision about which ports to block and which remain forwarding.

11.2.2 Cache Settings

Cache Size sets the maximum size of the bridge address cache, similar to the MAC or CAM table on a switch. The default is 100 entries. If there will be a large number of devices communicating across the bridge, set this higher.

Cache entry expire time controls the timeout of address cache entries in seconds. If set to 0, then address cache entries will not be expired. The default is 240 seconds (Four minutes).

11.2.3 Span Port

Selecting an interface as the Span port on the bridge will transmit a copy of every frame received by the bridge to the selected interface. This is most useful for snooping a bridged network passively on another host connected to the span ports of the bridge with something such as Snort, tcpdump, etc. The selected span port may not be a member port on the bridge.

11.2.4 Edge Ports / Automatic Edge Ports

If an interface is set as an Edge port, it is always assumed to be connected to an end device, and *never* to a switch; It assumes that the port can never create a layer 2 loop. Only set this on a port when it will never be connected to another switch. By default ports automatically detect edge status, and they can be selected under Auto Edge ports to *disable* this automatic edge detection behavior.

11.2.5 PTP Ports / Automatic PTP Ports

If an interface is set as a PTP port, it is always assumed to be connected to a switch, and not to an end user device; It assumes that the port can potentially create a layer 2 loop. It should only be enabled on ports that are connected to other RSTP-enabled switches. By default ports automatically detect PTP status, and they can be selected under Auto PTP ports to *disable* this automatic PTP detection behavior.

11.2.6 Sticky Ports

An interface selected in Sticky Ports will have its dynamically learned addresses cached as though they were static once they enter the cache. Sticky entries are never removed from the address cache, even if they appear on a different interface. This could be used a security measure to ensure that devices cannot move between ports arbitrarily.

11.2.7 Private Ports

An interface marked as a Private Port will not communicate with any other port marked as a Private Port. This can be used to isolate end users or sections of a network from each other if they are connected to separate bridge ports marked in this way. It works similar to “Private VLANs” or client isolation on a wireless access point.

11.3 Bridging and Interfaces

A bridge interface (e.g. *bridge0*) itself may be assigned as interface. This allows the bridge to act as a normal interface and have an IP address placed upon it rather than a member interface.

Configuring the IP address on the bridge itself is best in nearly all cases. The main reason for this is due to the fact that bridges are dependent on the state of the interface upon which the IP address is assigned. If the IP address for the bridge is configured on a member interface and that interface is down, the whole bridge will be down and no longer passing traffic. The most common case for this is a wireless interface bridged to an Ethernet LAN NIC. If the LAN NIC is unplugged, the wireless would be dead unless the IP address was configured on the bridge interface and not LAN. Another reason is that if limiters must be used for controlling traffic, then there must be an IP address on the bridge interface for them to work properly. Likewise, in order for Captive Portal or a transparent proxy to function on an internal bridge the IP address must be configured on the assigned bridge and not a member interface.

11.3.1 Swapping Interface Assignments

Before getting too far into talking about moving around bridge interface assignments, it must be noted that these changes should be made from a port that is not involved in the bridge. For example, if bridging WLAN to LAN, make the change from WAN or another OPT port. Alternately, download a backup of config.xml and manually make the changes. Attempting to make changes to a port while managing the firewall from that port will most likely result loss of access to the GUI, leaving the firewall unreachable.

Easy Method: Move settings to the new interface

The easiest, though not the quickest, path in the GUI is to remove the settings from the LAN interface individually (IP address, DHCP, etc) and then activate them on the newly assigned bridge interface.

Quick but Tricky: Reassign the Bridge as LAN

Though this method is a bit trickier than moving the settings, it can be much faster especially in cases where there are lots of firewall rules on LAN or a complex DHCP configuration. In this method, some hoop-jumping is required but ultimately the bridge ends up as the LAN interface, and it retains the LAN IP address, all of the former firewall rules, DHCP, and other interface configuration.

- Assign and configure the bridge members that have not yet been handled. Review the steps below to ensure the interface settings are correct even if the interfaces have already been assigned and configured.
  - Navigate to Interfaces > Assignments
  - Choose the interface from the Available network ports list
  - Click Add
  - Navigate to the new interface configuration page, e.g. Interfaces > OPT2
  - Check Enable
  - Enter a Description such as WiredLAN2
  - Set both IPv4 Configuration Type and IPv6 Configuration Type to *None*
  - Uncheck both Block private networks and Block bogon networks if checked
  - Click Save
  - Click Apply Changes
  - Repeat for additional unassigned future bridge members
- Create the new bridge
  - Navigate to Interfaces > Assignments on the Bridges tab
  - Click Add to create a new bridge
  - Enter a Description, such as LAN Bridge
  - Select all of the new bridge members EXCEPT the *LAN* interface in the Member interfaces list
  - Click Save
- Change the bridge filtering System Tunable to disable member interface filtering
  - Navigate to System > Advanced, System Tunables tab
  - Locate the entry for net.link.bridge.pfil_member or create a new entry if one does not exist, using that name for the Tunable
  - Click ⧉ to edit an existing entry
  - Enter 0 in the Value field
  - Click Save
- Navigate to Interfaces > Assignments
- Change the assignment of LAN to bridge0
- Click Save

- Assign and configure the old LAN interface as described previously, setting its IP configuration types to *None* and naming it WiredLAN
- Edit the bridge and select the newly assigned WiredLAN as a bridge member
- Change the bridge filtering System Tunable to enable bridge interface filtering
  - Use the procedure described previously, but set net.link.bridge.pfil_bridge to 1

Now the former LAN interface, along with the new bridge members, are all on a common layer 2 with the bridge assigned as LAN along with the other configuration.

Quickest but Most Difficult: Hand Edit config.xml

Hand editing config.xml can be very fast for those familiar with the configuration format in XML. This method is easy to get wrong, however, so be sure to have backups and install media nearby in case a mistake is made.

When hand editing config.xml to accomplish this task, do as follows:

- Assign the additional bridge members and set their IP configuration types to *None*
- Create the bridge, including *LAN* and *LAN2* and other bridge members
- Assign the bridge (e.g. as OPT2) and enable it, also with an IP configuration type of *None*
- Download a backup of config.xml from Diagnostics > Backup/Restore
- Open config.xml in a text editor that understands UNIX line endings
- Change the *LAN* assignment to bridge0
- Change the former *LAN* assignment to what used to be the bridge (e.g. *OPT2*)
- Edit the bridge definition to refer to *OPT2* and not *LAN*
- Save the changes
- Restore the edited config.xml from Diagnostics > Backup/Restore

The firewall will reboot with the desired setup. Monitor the console to ensure the settings were applied correctly and no errors are encountered during the boot sequence.

11.3.2 Assigned Bridge MAC Addresses and Windows

The MAC address for a bridge is determined randomly when the bridge is created, either at boot time or when a new bridge is created. That means that on each reboot, the MAC address can change. In many cases this does not matter, but Windows Vista, 7, 8, and 10 use the MAC address of the gateway to determine if they are on a specific network. If the MAC changes, the network identity will change and its status as public, private, etc. may need to be corrected. To work around this, enter a MAC address on the assigned bridge interface to spoof it. Then clients will always see the same MAC for the gateway IP address.

11.4 Bridging and firewalling

Filtering with bridged interfaces functions similar to routed interfaces, but there are some configuration choices to alter exactly how the filtering behaves. By default, firewall rules are applied on each member interface of the bridge on an inbound basis, like any other routed interface.

It is possible to decide whether the filtering happens on the bridge member interfaces, or on the bridge interface itself. This is controlled by two values on System > Advanced on the System Tunables tab, as seen in Figure *Bridge Filtering Tunables*. The net.link.bridge.pfil_member tunable controls whether or not the rules will be honored on the bridge member interfaces. By default, this is on (1). The net.link.bridge.pfil_bridge tunable controls whether or not the rules will be honored on the bridge interface itself. By default, this is off (0). At least one of these must be set to 1.

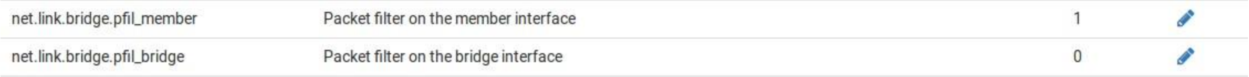

net.link.bridge.pfil_member	Packet filter on the member interface	1	✎
net.link.bridge.pfil_bridge	Packet filter on the bridge interface	0	✎

Fig. 1: Bridge Filtering Tunables

When filtering on the bridge interface itself, traffic will hit the rules as it enters from any member interface. The rules are still considered "inbound" like any other interface rules, but they work more like an interface group since the same rules apply to each member interface.

11.4.1 Firewall Rule Macros

Only one interface of a bridge will have an IP address set, the others will have none. For these interfaces, their firewall macros such as *OPT1 address* and *OPT1 net* are undefined because the interface has no address and thus no subnet.

If filtering is performed on bridge members, keep this fact in mind when crafting rules and explicitly list the subnet or use the macros for the interface where the IP address resides.

11.5 Bridging Two Internal Networks

When bridging two internal networks as described in *Internal Bridges* there are some special considerations to take for certain services on the firewall.

Note: There are additional requirements and restrictions when bridging wireless interfaces because of the way 802.11 functions. See *Bridging and wireless* for more information.

11.5.1 DHCP and Internal Bridges

When bridging one internal network to another, two things need to be done. First, ensure that DHCP is only running on the interface containing the IP address and not the bridge members without an address. Second, an additional firewall rule may be necessary at the top of the rules on the member interfaces to allow DHCP traffic.

Note: This only applies to filtering being performed on member interfaces, not filtering performed on the bridge.

When creating a rule to allow traffic on an interface, normally the source is specified similar to *OPT1 Subnet* so that only traffic from that subnet is allowed out of that segment. With DHCP, that is not enough. Because a client does not

16.4. Bridging and firewalling

yet have an IP address, a DHCP request is performed as a broadcast. To accommodate these requests, create a rule on the bridge member interfaces with the following settings:

- Navigate to Firewall > Rules on the tab for the bridge member
- Click ↴ Add to add a new rule to the top of the list
- Protocol: *UDP*
- Source: 0.0.0.0
- Source Port: 68
- Destination: 255.255.255.255
- Destination port: 67
- Description stating this will Allow DHCP
- Click Save and Apply Changes

The rule will look like Figure *Firewall rule to allow DHCP*.

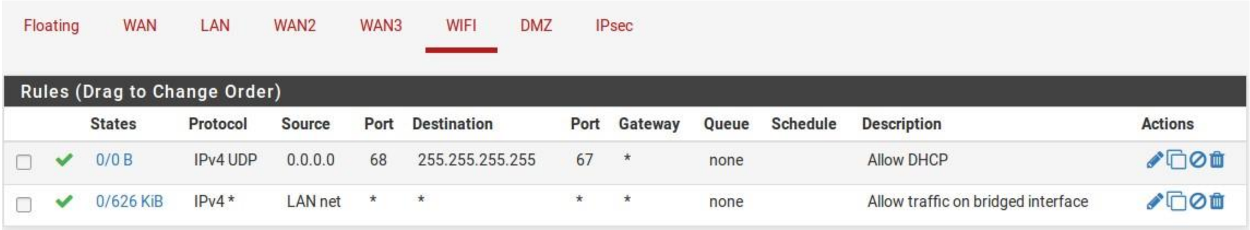

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
□	✓	0/0 B	IPv4 UDP	0.0.0.0	68	255.255.255.255	67	*	none		Allow DHCP	
□	✓	0/626 KiB	IPv4 *	LAN net	*	*	*	*	none		Allow traffic on bridged interface	

Fig. 2: Firewall rule to allow DHCP

After adding the rule, clients in the bridged segment will be able to successfully make requests to the DHCP daemon listening on the interface to which it is bridged.

DHCPv6 is a bit more complicated to allow since it communicates to and from both link-local and multicast IPv6 addresses. See Figure *Firewall Rule to Allow both DHCP and DHCPv6* for the list of required rules. These can be simplified with aliases into one or two rules containing the proper source network, destination network, and ports.

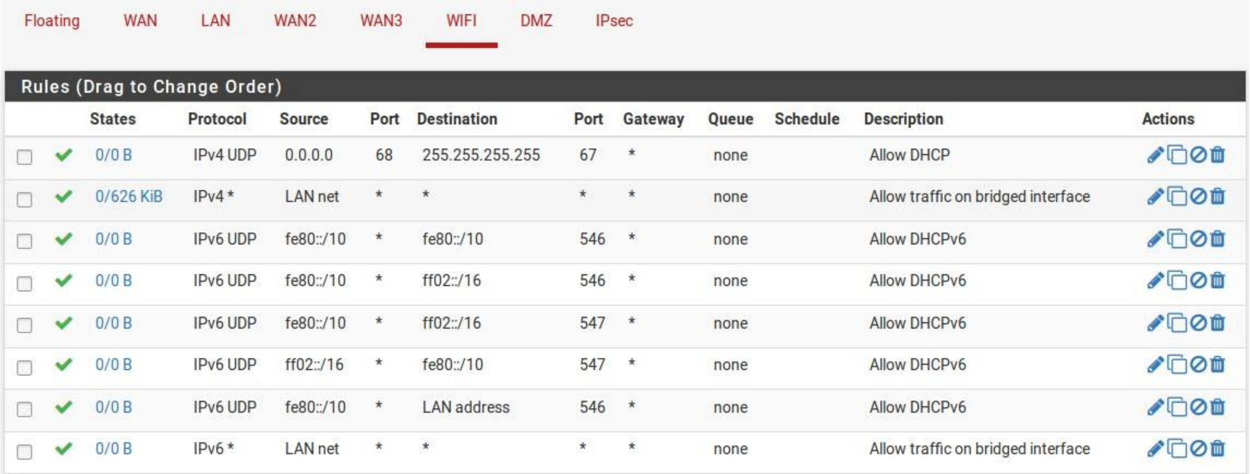

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
□	✓	0/0 B	IPv4 UDP	0.0.0.0	68	255.255.255.255	67	*	none		Allow DHCP	
□	✓	0/626 KiB	IPv4 *	LAN net	*	*	*	*	none		Allow traffic on bridged interface	
□	✓	0/0 B	IPv6 UDP	fe80::/10	*	fe80::/10	546	*	none		Allow DHCPv6	
□	✓	0/0 B	IPv6 UDP	fe80::/10	*	ff02::/16	546	*	none		Allow DHCPv6	
□	✓	0/0 B	IPv6 UDP	fe80::/10	*	ff02::/16	547	*	none		Allow DHCPv6	
□	✓	0/0 B	IPv6 UDP	ff02::/16	*	fe80::/10	547	*	none		Allow DHCPv6	
□	✓	0/0 B	IPv6 UDP	fe80::/10	*	LAN address	546	*	none		Allow DHCPv6	
□	✓	0/0 B	IPv6 *	LAN net	*	*	*	*	none		Allow traffic on bridged interface	

Fig. 3: Firewall Rule to Allow both DHCP and DHCPv6

11.6 Bridging interoperability

Bridged interfaces are different from normal interfaces, thus there are a few features that are incompatible with bridging and others where additional considerations are necessary to accommodate bridging. This section covers features that work differently with bridging than with non-bridged interfaces.

11.6.1 Captive portal

Captive portal (*Captive Portal*) is not compatible with transparent bridging because it requires an IP address on the interface being bridged, used to serve the portal contents, and that IP address must be the gateway for clients. This means that it is not possible, for example, to bridge LAN and WAN and hope to capture clients with the portal.

This can work when bridging multiple *local* interfaces to all route through AZTCO-FW® (e.g. LAN1, LAN2, LAN3, etc). It will work if the bridge interface is assigned, the bridge interface has an IP address, and that IP address is used as the gateway by clients on the bridge. See *Swapping Interface Assignments* for a procedure to place the IP address on an assigned bridge interface.

11.6.2 High Availability

High availability (*High Availability*) is not recommended with bridging. Some users have had mixed success with combining the two in the past but great care must be taken to handle layer 2 loops, which are unavoidable in a HA+bridge scenario. When two network segments are bridged, they are in effect merged into one larger network, as discussed earlier in this chapter. When HA is added into the mix, that means there will be two paths between the switches for each respective interface, creating a loop.

Managed switches can handle this with Spanning Tree Protocol (STP) but unmanaged switches have no defenses against looping. Left unchecked, a loop will bring a network to its knees and make it impossible to pass any traffic. STP may be configured on bridges to help, though there may still be unexpected results.

Consult switch documentation for information on STP configuration. Configure the switch to give preference to the port(s) on the primary node.

11.6.3 Multi-WAN

Transparent bridging by its nature is incompatible with multi-WAN in many of its uses. When using bridging between a WAN and LAN/OPT interface, commonly something other than AZTCO-FW will be the default gateway for the hosts on the bridged interface, and that router is the only device that can direct traffic from those hosts. This doesn’t prevent multi-WAN from being used with other interfaces on the same firewall that are not bridged, it only impacts the hosts on bridged interfaces where they use something other than AZTCO-FW as their default gateway. If multiple internal interfaces are bridged together and AZTCO-FW is the default gateway for the hosts on the bridged interfaces, then multi-WAN can be used the same as with non-bridged interfaces.

11.6.4 Limiters

For limiters to function with bridging, the bridge itself must be assigned and the bridge interface must have the IP address and not a member interface.

11.6.5 LAN NAT and Transparent Proxies

For port forwards on LAN, or transparent proxies which use port forwards on LAN to capture traffic, to function in a bridge scenario, the situation is the same as Captive Portal: It will only function for LAN bridges and not WAN/LAN bridges, the IP address must be on the assigned bridge interface, and that IP address must be used as the gateway for local clients.

This means that a package such as Squid cannot work in a transparent firewall scenario where LAN is bridged to a WAN.

11.6.6 Mixing Bridged and NAT Segments

For hosts behind the NAT/routed segment, NAT must occur as traffic exits toward the bridged systems so that the return traffic will come back to the firewall.

For hosts on the bridged segment to reach hosts behind the NAT segment directly, a static route could be used on the bridged hosts or upstream gateway to send the “private” subnet traffic to the IP address of the firewall in the bridged network.

Normally each interface on the AZTCO-FW® firewall represents its own broadcast domain with a unique IP subnet. In some circumstances it is desirable or necessary to combine multiple interfaces onto a single broadcast domain, where two ports on the firewall will act as if they are on the same switch, except traffic between the interfaces can be controlled with firewall rules. Typically this is done so multiple interfaces will act as though they are on the same flat network using the same IP subnet and so that clients all share broadcast and multicast traffic.

Certain applications and devices rely on broadcasts to function, but these are found more commonly in home environments than corporate environments. For a practical discussion, see *Bridging and wireless*.

For services running on the firewall, bridging can be problematic. Features such as limiters, Captive Portal, and transparent proxies require special configuration and handling to work on bridged networks. Specifically, the bridge itself must be assigned and the only interface on the bridge with an IP address must be the assigned bridge. Also, in order for these functions to work, the IP address on the bridge must be the address used by clients as their gateway. These issues are discussed more in-depth in *Bridging interoperability*.

11.7 Types of Bridges

There are two distinct types of bridges: Internal bridges and Internal/external bridges. Internal bridges connect two local interfaces such as two LAN interfaces or a LAN interface and a wireless interface. Internal/external bridges connect a LAN to a WAN resulting in what is commonly called a “transparent firewall”.

16.7. Types of Bridges

11.7.1 Internal Bridges

With an internal type bridge, ports on the firewall are linked such that they behave similar to switch ports, though with the ability to filter traffic on the ports or bridge and with much lower performance than a switch. The firewall itself is still visible to the local connected clients and acts as their gateway, and perhaps DNS and DHCP server. Clients on the bridged segments may not even know there is a firewall between them.

This type of configuration is commonly chosen by administrators to isolate and control a portion of the network, such as a wireless segment, or to make use of additional ports on the firewall in lieu of a proper switch where installing a switch would be impractical. Though it is not recommended, this type of bridge can also be used to join two remote networks over certain types of VPN connections.

11.7.2 Internal/External Bridges

An Internal/External type bridge, also known as a "transparent firewall", is used to insert a firewall between two segments without altering the other devices. Most commonly this is used to bridge a WAN to an internal network so that the WAN subnet may be used "inside" the firewall, or internally between local segments as an in-line filter. Another common use is for devices behind the firewall to obtain IP addresses via DHCP from an upstream server on the WAN.

In a transparent firewall configuration the firewall does not receive the traffic directly or act as a gateway, it merely inspects the traffic as it passes through the firewall.

Note: Devices on the internal side of this bridge must continue to use the upstream gateway as their own gateway. Do not set any IP address on the firewall as a gateway for devices on a transparent bridge.

NAT is not possible with this style of bridge because NAT requires the traffic to be addressed to the firewall's MAC address directly in order to take effect. Since the firewall is not the gateway, this does not happen. As such, rules to capture traffic such as those used by a transparent proxy do not function.

11.8 Bridging and Layer 2 Loops

When bridging, care must be taken to avoid layer 2 loops, or a switch configuration must be in place that handles loops. A layer 2 loop is when, either directly or indirectly, the switch has a connection back to itself. If a firewall running AZTCO-FW has interfaces bridged together, and two interfaces are plugged into the same switch on the same VLAN, a layer 2 loop has been created. Connecting two patch cables between two switches also does this.

Managed switches employ Spanning Tree Protocol (STP) to handle situations like this, because it is often desirable to have multiple links between switches, and the network shouldn't be exposed to complete meltdown by someone plugging one network port into another network port. STP is not enabled by default on all managed switches, and is almost never available with unmanaged switches. Without STP, the result of a layer 2 loop is frames on the network will circle endlessly and the network will completely cease to function until the loop is removed. Check the switch configuration to ensure the feature is enabled and properly configured.

AZTCO-FW enables STP on bridge interfaces to help with loops, but it can still lead to unexpected situations. For instance, one of the bridge ports would shut itself down to stop the loop, which could cause traffic to stop flowing unexpectedly or bypass the firewall entirely.

16.8. Bridging and Layer 2 Loops

In a nutshell, bridging has the potential to completely melt down the network unless anyone that plugs devices into the switch is careful.

CHAPTER

TWELVE

VIRTUAL LANS (VLANS)

12.1 Terminology

This section defines the terminology required to successfully deploy VLANs.

Trunking Trunking refers to a means of carrying multiple VLANs on the same physical switch port. The frames leaving a trunk port are marked with an 802.1Q tag in the header, enabling the connected device to differentiate between multiple VLANs. Trunk ports are used to connect multiple switches, and for connecting any devices that are capable of 802.1Q tagging and require access to multiple VLANs. This is commonly limited to the firewall or router providing connectivity between VLANs, in this case, the AZTCO-FW® firewall, as well as any connections to other switches containing multiple VLANs.

VLAN ID Each VLAN has an identifier number (ID) for distinguishing tagged traffic. This is a number between *1* and *4094*. The default VLAN on switches is VLAN *1*, and this VLAN should not be used when deploying VLAN trunking. This is discussed further in *VLANs and Security*. Aside from avoiding the use of VLAN 1, VLAN numbers may be chosen at will. Some designs start with VLAN 2 and increment by one until the required number of VLANs is reached. Another common design is to use the third octet in the subnet of the VLAN as the VLAN ID. For example, if 10.0.10.0/24, 10.0.20.0/24 and 10.0.30.0/24 are used, it is logical to use VLANs 10, 20, and 30 respectively. Choose a VLAN ID assignment scheme that makes sense for a given network design.

Parent interface The physical interface where a VLAN resides is known as its Parent Interface. For example, *igb0* or *em0*. When VLANs are configured on AZTCO-FW, each is assigned a virtual interface. The virtual interface name is crafted by combining the parent interface name plus the VLAN ID. For example, for VLAN 20 on *igb0*, the interface name is igb0_vlan20.

Note: The sole function of the parent interface is, ideally, to be the parent for the defined VLANs and not used directly. In some situations this will work, but can cause difficulties with switch configuration, and it requires use of the default VLAN on the trunk port, which is best to avoid as discussed further in *VLANs and Security*.

Access Port An access port refers to a switch port providing access to a single VLAN, where the frames are not tagged with an 802.1Q header. Normal client-type devices are connected to access ports, which will comprise the majority of switch ports. Devices on access ports do not need knowledge of VLANs or tagging. They see the network on their port the same as they would a switch without VLANs.

Double tagging (QinQ) *QinQ* refers to the double tagging of traffic, using both an outer and inner 802.1Q tag. This can be useful in large ISP environments, other very large networks, or networks that must carry multiple VLANs across a link that only supports a single VLAN tag. Triple tagging is also possible. AZTCO-FW supports QinQ, though it is not a very commonly used feature. These types

of environments generally need the kind of routing power that only a high end ASIC-based router can support, and QinQ adds a level of complexity that is unnecessary in most environments. For more information on configuring QinQ on AZTCO-FW, see *AZTCO-FW QinQ Configuration*.

Private VLAN (PVLAN) PVLAN, sometimes called Port Isolation, refers to capabilities of some switches to segment hosts within a single VLAN. Normally hosts within a single VLAN function the same as hosts on a single switch without VLANs configured. PVLAN provides a means of preventing hosts on a VLAN from talking to any other host on that VLAN, only permitting communication between that host and its default gateway. This isn't directly relevant to AZTCO-FW, but is a common question. Switch functionality such as this is the only way to prevent communication between hosts in the same subnet. Without a function like PVLAN, no network firewall can control traffic within a subnet because it never touches the default gateway.

12.2 VLANs and Security

VLANs are a great way to segment a network and isolate subnetworks, but there are security issues which need to be taken into account when designing and implementing a solution involving VLANs. VLANs are not inherently insecure, but misconfiguration can leave a network vulnerable. There have also been past security problems in switch vendor implementations of VLANs.

12.2.1 Segregating Trust Zones

Because of the possibility of misconfiguration, networks of considerably different trust levels should be on separate physical switches. For example, while the same switch could technically be used with VLANs for all internal networks as well as the network outside the firewalls, that should be avoided as a simple misconfiguration of the switch could lead to unfiltered Internet traffic entering the internal network. At a minimum, use two switches in such scenarios: One for outside the firewall and one inside the firewall. In many environments, DMZ segments are also treated separately, on a third switch in addition to the WAN and LAN switches. In others, the WAN side is on its own switch, while all the networks behind the firewall are on the same switches using VLANs. Which scenario is most appropriate for a given network depends on its specific circumstances, level of risk, and security concerns.

12.2.2 Using the default VLAN1

Because VLAN 1 is the default ("native") VLAN, it may be used in unexpected ways by the switch. It is similar to using a default-allow policy on firewall rules instead of default deny and selecting what is needed. Using a different VLAN is always better, and ensure that only the ports are selected that must be on that VLAN, to better limit access. Switches will send internal protocols such as STP (Spanning Tree Protocol), VTP (VLAN Trunking Protocol), and CDP (Cisco Discover Protocol) untagged over the native VLAN, where the switches use these protocols. It is generally best to keep that internal traffic isolated from data traffic.

If VLAN 1 must be used, take great care to assign every single port on every switch to a different VLAN except those that must be in VLAN 1, and do not create a management interface for the switch on VLAN 1. The native VLAN of the switch group should also be changed to a different, unused, VLAN. Some switches may not support any of these workarounds, and so it is typically easier to move data to a different VLAN instead of fussing with making VLAN 1 available. With VLAN ID 2 through 4094 to choose from, it is undoubtedly better to ignore VLAN 1 when designing a new VLAN scheme.

12.2.3 Using a trunk port's default VLAN

When VLAN tagged traffic is sent over a trunk on the native VLAN, tags in the packets that match the native VLAN may be stripped by the switch to preserve compatibility with older networks. Worse yet, packets that are double tagged with the native VLAN and a different VLAN will only have the native VLAN tag removed when trunking in this way and when processed later, that traffic can end up on a different VLAN. This is also called "VLAN hopping".

As mentioned in the previous section, any untagged traffic on a trunk port will be assumed to be the native VLAN, which could also overlap with an assigned VLAN interface. Depending on how the switch handles such traffic and how it is seen by AZTCO-FW®, using the interface directly could lead to two interfaces being on the same VLAN.

12.2.4 Limiting access to trunk ports

Because a trunk port can talk to any VLAN in a group of trunked switches, possibly even ones not present on the current switch depending on the switch configurations, it is important to physically secure trunk ports. Also make sure there are no ports configured for trunking that are left unplugged and enabled where someone could hook into one, accidentally or otherwise. Depending on the switch, it may support dynamic negotiation of trunking. Ensure this functionality is disabled or properly restricted.

12.2.5 Other Issues with Switches

Over the years there have been reports of rare cases where VLAN-based switches have leaked traffic across VLANs while under heavy loads, or if a MAC address of a PC on one VLAN is seen on another VLAN. These issues tend to be in older switches with outdated firmware, or extremely low-quality managed switches. These types of issues were largely resolved many years ago, when such security problems were common. No matter what switch from what brand is used for a network, research to see if it has undergone any kind of security testing, and ensure the latest firmware is loaded on the switch. While these issues are a problem with the switch, and not AZTCO-FW, they are part of a network's overall security.

Many of the items here are specific to particular makes and models of switches. Security considerations differ based on the switch being used on a network. Refer to its documentation for recommendations on VLAN security.

12.3 AZTCO-FW VLAN Configuration

This section covers how to configure VLANs in AZTCO-FW® software.

12.3.1 Console VLAN configuration

VLANs can be configured at the console using the Assign Interfaces function. The following example shows how to configure two VLANs, ID 10 and 20, with igb2 as the parent interface. The VLAN interfaces are assigned as OPT1 and OPT2:

```
0) Logout (SSH only)                      8) pfTop
1) Assign Interfaces                     9) Filter Logs
2) Set interface(s) IP address           10) Restart webConfigurator
3) Reset webConfigurator password        11) AZTCO-FW  Shell
4) Reset to factory defaults             12) Update from console
5) Reboot system                         13) Disable Secure Shell (sshd)
6) Halt system                           14) Restore recent configuration
7)    Ping host                          15) Restart PHP-FPM
```

(continues on next page)

(continued from previous

```
Enter an option: 1 Valid

interfaces are:

igb0        00:08:a2:09:95:b5  (up) Intel(R) PRO/1000 Network Connection, Version igb1  00:08:a2:09:95:b6  (up) Intel(R)
PRO/1000 Network Connection, Version igb2      00:08:a2:09:95:b1 (down) Intel(R) PRO/1000 Network Connection,
Version igb3          00:08:a2:09:95:b2 (down) Intel(R) PRO/1000 Network Connection, Version igb4
            00:08:a2:09:95:b3 (down) Intel(R) PRO/1000 Network Connection, Version igb5            00:08:a2:09:95:b3
(down) Intel(R) PRO/1000 Network Connection, Version -

Do VLANs need to be set up first?
If VLANs will not be used, or only for optional interfaces, it is typical to say no here and use the webConfigurator to
configure VLANs later, if required.

Should VLANs be set up now [y|n]? y

WARNING: all existing VLANs will be cleared if you proceed!

Do you want to proceed [y|n]? y VLAN Capable

interfaces:

igb0        00:08:a2:09:95:b5  (up) igb1
            00:08:a2:09:95:b6  (up) igb2
            00:08:a2:09:95:b1 igb3
            00:08:a2:09:95:b2 igb4
            00:08:a2:09:95:b3  (up) igb5
            00:08:a2:09:95:b3  (up)

Enter the parent interface name for the new VLAN (or nothing if finished): igb2
Enter the VLAN tag (1-4094): 10 VLAN Capable

interfaces:

igb0        00:08:a2:09:95:b5  (up) igb1
            00:08:a2:09:95:b6  (up) igb2
            00:08:a2:09:95:b1 igb3
            00:08:a2:09:95:b2 igb4
            00:08:a2:09:95:b3  (up) igb5
            00:08:a2:09:95:b3  (up)

Enter the parent interface name for the new VLAN (or nothing if finished): igb2
Enter the VLAN tag (1-4094): 20 VLAN Capable

interfaces:

igb0        00:08:a2:09:95:b5  (up) igb1
            00:08:a2:09:95:b6  (up) igb2
            00:08:a2:09:95:b1 igb3
            00:08:a2:09:95:b2 igb4
```

```
        00:08:a2:09:95:b3   (up) igb5
        00:08:a2:09:95:b3   (up)

Enter the parent interface name for the new VLAN (or nothing if finished): <enter>
```

(continues on next page) (continued from previous page)

```
VLAN interfaces:

igb2_vlan10        VLAN tag 10, parent interface igb2 igb2_vlan20    VLAN tag 20, parent
interface igb2

If the names of the interfaces are not known, auto-detection can be used instead. To use auto-
detection, please disconnect all interfaces before pressing 'a' to begin the process.

Enter the WAN interface name or 'a' for auto-detection
(igb0 igb1 igb2 igb3 igb4 igb5 igb2_vlan10 igb2_vlan20 or a): igb1

Enter the LAN interface name or 'a' for auto-detection NOTE: this enables full
Firewalling/NAT mode.
(igb0 igb2 igb3 igb4 igb5 igb2_vlan10 igb2_vlan20 a or nothing if finished): igb0

Enter the Optional 1 interface name or 'a' for auto-detection
(igb2 igb3 igb4 igb5 igb2_vlan10 igb2_vlan20 a or nothing if finished): igb2_vlan10

Enter the Optional 2 interface name or 'a' for auto-detection
(igb2 igb3 igb4 igb5 igb2_vlan20 a or nothing if finished): igb2_vlan20

Enter the Optional 3 interface name or 'a' for auto-detection
(igb2 igb3 igb4 igb5 a or nothing if finished):<enter> The interfaces will be assigned

as follows:

WAN -> igb1
LAN -> igb0
OPT1 -> igb2_vlan10
OPT2 -> igb2_vlan20

Do you want to proceed [y|n]? y

Writing configuration...done.
One moment while the settings are reloading... done!
```

After a few seconds, the firewall settings will reload and the console menu will reload.

12.3.2 Web interface VLAN configuration

In the system used for this example, WAN and LAN are assigned as *igb1* and *igb0* respectively. There is also an *igb2* interface that will be used as the VLAN parent interface.

To configure VLANs in the AZTCO-FW web interface:

- Navigate to Interfaces > Assignments to view the interface list.
- Click the VLANs tab.
- Click + Add to add a new VLAN
- Configure the VLAN as shown in Figure *Edit VLAN*.

Parent Interface The physical interface upon which this VLAN tag will be used. In this case, *igb2*

VLAN tag The VLAN ID number, in this case, 10

VLAN Priority Leave at the default value, blank

Description Some text to identify the purpose of the VLAN, such as DMZ

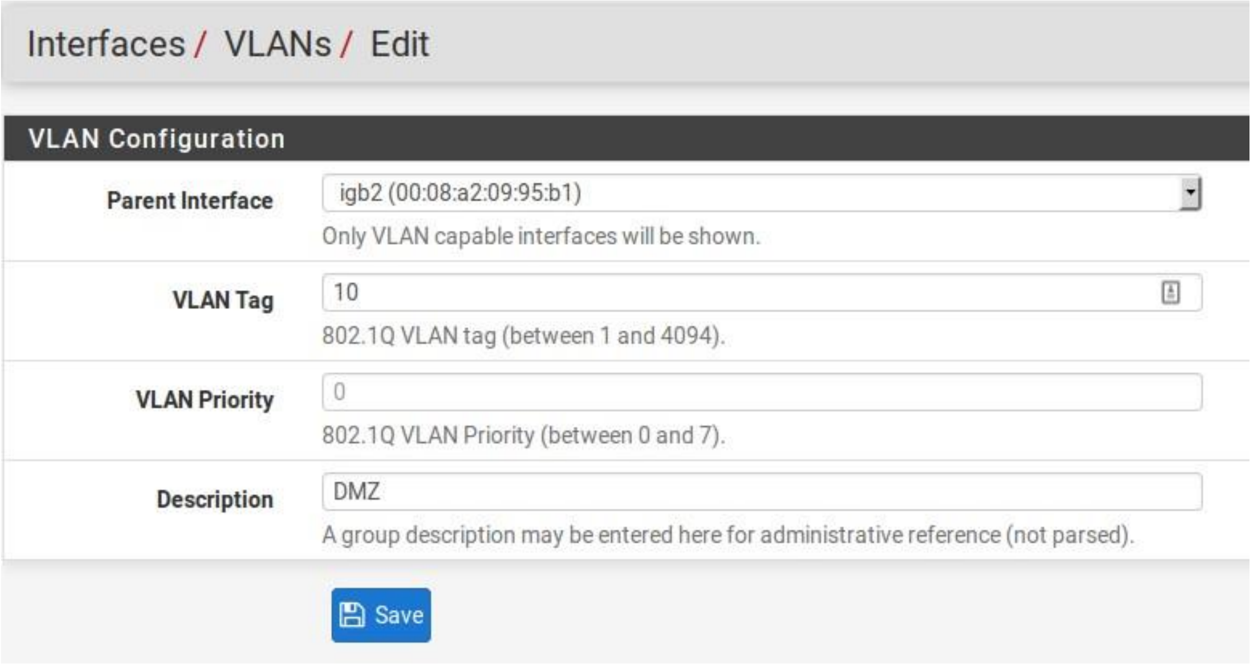

Interfaces / VLANs / Edit

VLAN Configuration

Parent Interface	igb2 (00:08:a2:09:95:b1) Only VLAN capable interfaces will be shown.
VLAN Tag	10 802.1Q VLAN tag (between 1 and 4094).
VLAN Priority	0 802.1Q VLAN Priority (between 0 and 7).
Description	DMZ A group description may be entered here for administrative reference (not parsed).

Fig. 1: Edit VLAN

- Click Save to return to the VLAN list, which now includes the newly added VLAN 10.
- Repeat the process to add additional VLANs, such as VLAN 20. These can be seen in Figure *VLAN list*

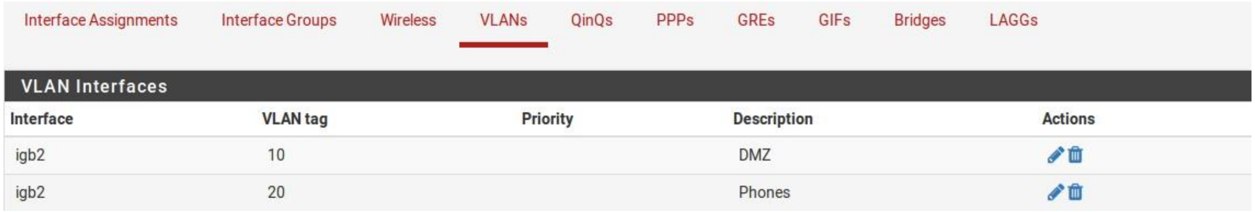

VLAN Interfaces

Interface	VLAN tag	Priority	Description	Actions
igb2	10		DMZ	✎ 🗑
igb2	20		Phones	✎ 🗑

Fig. 2: VLAN list

To assign the VLANs to interfaces:

- Navigate to Interfaces > Assignments
- Click the Interface Assignments tab
- Select the VLAN to add from the Available Network Ports list, such as *VLAN 10 on igb2 (DMZ)*
- Click + Add to assign the network port
- Repeat the last two steps to assign *VLAN 20 on igb2 (Phones)*

When finished, the interfaces will look like Figure *Interfaces list with VLANs*

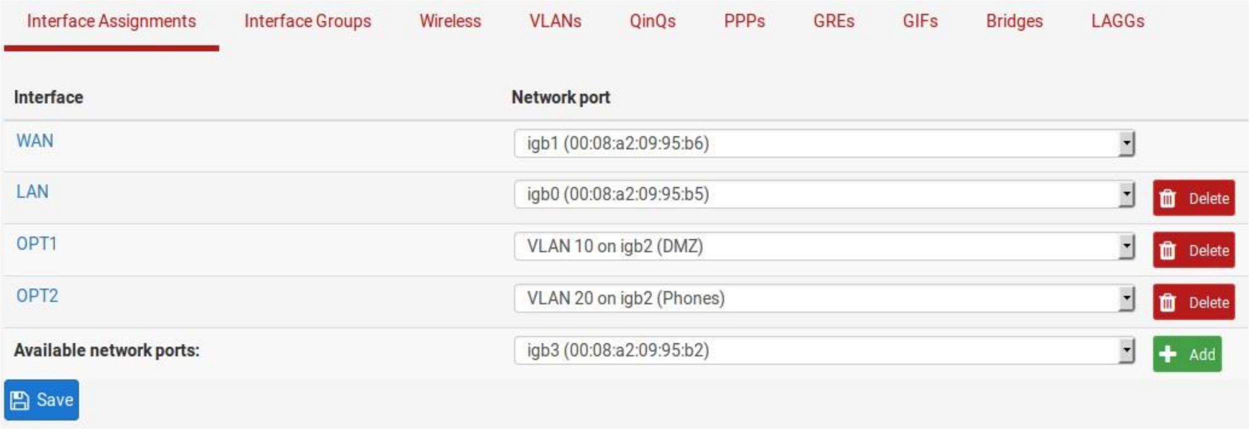

Interface	Network port
WAN	igb1 (00:08:a2:09:95:b6)
LAN	igb0 (00:08:a2:09:95:b5)
OPT1	VLAN 10 on igb2 (DMZ)
OPT2	VLAN 20 on igb2 (Phones)
Available network ports:	igb3 (00:08:a2:09:95:b2)

Fig. 3: Interfaces list with VLANs

The VLAN-based OPT interfaces behave as any other OPT interfaces do, which means they must be enabled, configured, have firewall rules added, and services like the DHCP Server will need to be configured if needed. See *Interface Configuration Basics* for more information on configuring optional interfaces.

See also:

- *Configuring Switches with VLANs*
- *AZTCO-FW QinQ Configuration*

VLANs enable a switch to carry multiple discrete broadcast domains, allowing a single switch to function as if it were multiple switches. VLANs are commonly used for network segmentation in the same way that multiple switches can be used: To place hosts on a specific segment, isolated from other segments. Where trunking is employed between switches, devices on the same segment need not reside on the same switch. Devices that support trunking can also communicate on multiple VLANs through a single physical port.

This chapter covers VLAN concepts, terminology and configuration in AZTCO-FW® software.

12.4 Requirements

There are two requirements, both of which must be met to deploy VLANs.

1. 802.1Q VLAN capable switch

   Every decent managed switch manufactured in the last 15 years supports 802.1Q VLAN trunking.

   Warning: VLANs cannot be used with an unmanaged switch.

2. Network adapter capable of VLAN tagging

   A NIC that supports hardware VLAN tagging or has long frame support is required. Each VLAN frame has a 4 byte 802.1Q tag added in the header, so the frame size can be up to 1522 bytes. A NIC supporting hardware VLAN tagging or long frames is required because other adapters will not function with frames larger than the normal 1518 byte maximum with 1500 MTU Ethernet. This will cause large frames to be dropped, which causes performance problems and connection stalling.

Note: If an adapter is listed as having long frame support does not guarantee the specific implementation of

17.4. Requirements

that NIC chipset properly supports long frames. Realtek *rl(4)* NICs are the biggest offenders. Many will work fine, but some do not properly support long frames, and some will not accept 802.1Q tagged frames at all. If problems are encountered using one of the NICs listed under long frame support, we recommend trying an interface with VLAN hardware tagging support instead. We are not aware of any similar problems with NICs listed under VLAN hardware support.

Ethernet interfaces with VLAN hardware support:

ae(4), *age(4)*, *alc(4)*, *ale(4)*, *bce(4)*, *bge(4)*, *bxe(4)*, *cxgb(4)*, *cxgbe(4)*, *em(4)*, *igb(4)*, *ixgb(4)*, *ixgbe(4)*, *jme(4)*, *msk(4)*, *mxge(4)*, *nxge(4)*, *nge(4)*, *re(4)*, *sge(4)*, *stge(4)*, *ti(4)*, *txp(4)*, *vge(4)*.

Ethernet interfaces with long frame support : *axe(4)*, *bfe(4)*, *cas(4)*, *dc(4)*, *et(4)*, *fwe(4)*, *fxp(4)*, *gem(4)*,

hme(4), *le(4)*, *nfe(4)*, *nve(4)*, *rl(4)*, *sf(4)*, *sis(4)*, *sk(4)*, *ste(4)*, *tl(4)*, *tx(4)*, *vr(4)*, *vte(4)*, *xl(4)*.

CHAPTER

THIRTEEN

MULTIPLE WAN CONNECTIONS

13.1 Multi-WAN Terminology and Concepts

This section covers the terminology and concepts necessary to understand to deploy multi-WAN with AZTCO-FW® software.

13.1.1 WAN-type Interface

A WAN-type interface is an interface through which the Internet can be reached, directly or indirectly. The firewall treats any interface with a gateway selected on its Interfaces menu page as a WAN. For example, with a static IP address WAN, Interfaces > WAN has a gateway selected, such as WAN_GW. If this gateway selection is not present, then the interface will be treated as a local interface instead. Dynamic IP address interfaces such as DHCP and PPPoE receive a dynamic gateway automatically and are always treated as WANs.

The presence of a gateway on the interface configuration changes the firewall behavior on such interfaces in several ways:

- Firewall rules on these interfaces have reply-to added which returns connections coming in through that WAN back out via the same WAN where possible
- These interfaces are used as exit interfaces for automatic and hybrid outbound NAT
- These interfaces are treated as WANs by the traffic shaper wizard

Warning: Do not select a gateway on the Interfaces menu entry for local interfaces such as LAN.

Local and other interfaces may have a gateway defined under System > Routing, so long as that gateway is not chosen under their interface configuration, for example on Interfaces > LAN.

13.1.2 Policy routing

Policy routing refers to a means of routing traffic by more than the destination IP address of the traffic, as is done with the routing table in most operating systems and routers. This is accomplished by the use of a policy of some sort, usually firewall rules or an access control list. In AZTCO-FW, the Gateway field available when editing or adding firewall rules enables the use of policy routing. The Gateway field contains all gateways defined on the firewall under System > Routing, plus any gateway groups.

Policy routing provides a powerful means of directing traffic to the appropriate WAN interface or other gateway, since it allows matching anything a firewall rule can match. Specific hosts, subnets, protocols and more can be used to direct traffic.

Note: Remember that all firewall rules, including policy routing rules, are processed in top down order, and the first match wins.

13.1.3 Gateway Groups

Gateway groups define how a chosen set of gateways provide failover and/or load balancing functionality. They are configured under System > Routing, on the Gateway Groups tab.

See *Gateway Groups* for more.

13.1.4 Failover

Failover refers to the ability to switch from one or more WANs to an alternate WAN if the preferred connection fails. This is useful for situations where traffic should utilize one specific WAN connection unless it is unavailable.

See also:

To fail from one *firewall* to another, rather than from one *WAN* to another, see *High Availability*.

13.1.5 Load Balancing

The Load Balancing functionality in AZTCO-FW software distributes connections over multiple WAN connections in a round-robin fashion. This feature operates on a per-connection basis. If a gateway that is part of a load balancing group fails, the interface is marked as down and removed from all groups until it recovers, thus a load balanced configuration effectively also includes failover functionality.

18.1.6 Monitor IP Addresses

When configuring failover or load balancing, each gateway is associated with a monitor IP address (*Gateway Settings*). In a typical configuration, the firewall will ping this IP address and if it stops responding, the gateway is marked as down. Options on the gateway group can select different failure triggers besides packet loss. The other triggers are high latency, a combination of either packet loss or high latency, or when the circuit is down.

What constitutes failure?

The topic is a little more complex than "if pings to the monitor IP address fail, the gateway is marked as down." The actual criteria for a failure depend on the options chosen when creating the gateway group and the individual settings on a gateway.

The settings for each gateway that control when it is considered up and down are all discussed in *Advanced Gateway Settings*. The thresholds for packet loss, latency, down time, and even the probing interval of the gateway are all individually configurable.

13.1.7 State Killing/Forced Switch

When a gateway has failed, the firewall can optionally flush all states to force clients to reconnect, and in doing so they will use a gateway that is online instead of a gateway that is down. This currently only works one-way, meaning that it can move connections off of a failing gateway, but it cannot force them back if the original gateway comes back online.

This is an optional behavior, but it is not enabled by default since it is disruptive. For information on changing this setting, see *Gateway Monitoring*.

13.1.8 Default Gateway Switching

Traffic exiting the firewall itself will use the default gateway unless a static route sends the packet along a different path. If the default gateway is on a WAN that is down, daemons on the firewall will be unable to make outbound connections, depending on the capabilities of the daemon and its configuration.

The default gateway for the firewall can be set to a gateway group or set to an automatic mode, which will switch the default to the next available gateway if the normal default gateway fails, and then switched back when that WAN recovers. See *Managing the Default Gateway* for details.

13.2 Policy Routing, Load Balancing and Failover Strategies

This section provides guidance on common Multi-WAN goals and how they are achieved with AZTCO-FW® software.

13.2.1 Bandwidth Aggregation

One of the primary desires with multi-WAN is bandwidth aggregation. Load balancing can help accomplish this goal. There is, however, one caveat: If the firewall has two 50 Mbps WAN circuits, it cannot get 100 Mbps of throughput with a *single* client connection. Each individual connection must be tied to only one specific WAN. This is true of any multi-WAN solution other than MLPPP. The bandwidth of two different Internet connections cannot be aggregated into a single large “pipe” without involvement from the ISP. With load balancing, since individual connections are balanced in a round-robin fashion, 100 Mbps of throughput can only be achieved using two 50 Mbps circuits when multiple connections are involved. Applications that utilize multiple connections, such as download accelerators, will be able to achieve the combined throughput capacity of the two or more connections.

Note: Multi-Link PPPoE (MLPPP) is the only WAN type which can achieve full aggregate bandwidth of all circuits in a bundle, but requires special support from the ISP. For more on MLPPP, see *Multi-Link PPPoE (MLPPP)*

In networks with numerous internal machines accessing the Internet, load balancing will reach speeds near the aggregate throughput by balancing the many internal connections out all of the WAN interfaces.

13.2.2 Segregation of Priority Services

Consider a site which has a reliable, high quality Internet connection that offers low bandwidth, or high costs for excessive transfers, and another connection that is fast but of lesser quality (higher latency, more jitter, or less reliable). In these situations, services can be segregated between the two Internet connections by their priority. High priority services may include VoIP, traffic destined to a specific network such as an outsourced application provider, or specific protocols used by critical applications, amongst other options. Low priority traffic commonly includes any permitted traffic that doesn't match the list of high priority traffic. Policy routing rules can be setup to direct the high priority traffic out the high quality Internet connection, and the lower priority traffic out the lesser quality connection.

Another example of a similar scenario is getting a dedicated Internet connection for quality critical services such as VoIP, and only using that connection for those services.

13.2.3 Failover Only

There are scenarios where only using failover is the best practice. For example, users who have a secondary backup Internet connection with a low bandwidth cap such as a 4G/LTE modem, and only want to use that connection if their primary connection fails, Gateway groups configured for failover can achieve this goal.

Another usage for failover is to ensure a certain protocol or destination always uses only one WAN unless it goes down.

18.2.4 Unequal Cost Load Balancing

AZTCO-FW software can achieve unequal cost load balancing by setting appropriate weights on the gateways as discussed in *Advanced Gateway Settings*. By setting a weight on a gateway, it will be used more often in a gateway group. Weights can be set from *1* to *30*, allowing

Table 1: Unequal Cost Load Balancing

WAN_GW weight	WAN2_GW weight	WAN load	WAN2 load
3	2	60%	40%
2	1	67%	33%
3	1	75%	25%
4	1	80%	20%
5	1	83%	17%
30	1	97%	3%

Note that this distribution is strictly balancing the number of *connections*, it does not take interface throughput into account. This means bandwidth usage will not necessary be distributed equally, though in most environments it works out to be roughly distributed as configured over time. This also means if an interface is loaded to its capacity with a single high throughput connection, additional connections will still be directed to that interface.

13.3 Multi-WAN Caveats and Considerations

This section contains the caveats and considerations specific to multi-WAN in AZTCO-FW® software.

13.3.1 Multiple WANs sharing a single gateway IP

Due to the way pf handles multi-WAN connections, traffic can only be directed using the gateway IP address of a circuit, which is fine for most scenarios. If the firewall has multiple connections on the same ISP using the same subnet and gateway IP address, as is common when using multiple cable modems, an intermediate NAT device must be used on all but one of them so that the firewall sees each WAN gateway as a unique IP address.

When using the NAT device, it can be configured to forward all traffic back to the firewall which can help with using that WAN for other services. However, some protocols, such as VoIP, will have problem if they use a WAN with NAT in such a configuration.

If at all possible, contact the ISP and have them configure the WAN circuits such that they are in different subnets with different gateways.

One exception to this is a PPP type WAN such as PPPoE. PPP type WANs are capable of having the same gateway on multiple interfaces, but each gateway entry must be configured to use a different monitor IP address (See *Gateway Settings*).

13.3.2 Multiple PPPoE WANs

When multiple PPPoE lines from the same ISP are present and the ISP supports Multi-Link PPPoE (MLPPP), it may be possible to bond the lines into a single aggregate link. This bonded link has total bandwidth of all lines together in a single WAN as seen by the firewall. Configuration of MLPPP is covered in *Multi-Link PPPoE (MLPPP)*.

13.3.3 Local Services and Multi-WAN

There are additional considerations with local services and multi-WAN, since any traffic initiated from the firewall itself will not be affected by policy routing configured on internal interface rules. Traffic from the firewall itself always follows the routing table. Hence static routes are required under some circumstances when using additional WAN interfaces, otherwise only the WAN interface with the default gateway would be used.

The firewall can be configured to change the default gateway if the preferred default fails. See *Managing the Default Gateway* for details.

In the case of traffic initiated on the Internet destined for any WAN interface, AZTCO-FW automatically uses the reply-to directive in all WAN-type interface rules, which ensures the reply traffic is routed back out the correct WAN interface.

DNS Resolver

The default settings for the DNS Resolver require using failover for the default gateway to work properly with MultiWAN. See *Managing the Default Gateway* for details. As an alternative to using default gateway switching, a few changes can be made to make the DNS Resolver more accommodating to Multi-WAN, including enabling forwarding mode. The details are described later in this chapter.

DNS Forwarder

The DNS servers used by the DNS forwarder must have gateways defined if they use an non-default WAN interface, as described later in this chapter. There are no other caveats to DNS forwarder in multi-WAN environments.

Dynamic DNS

Dynamic DNS entries can be set using a gateway group for their interface. This will move a Dynamic DNS entry between WANs in failover mode, allowing a public hostname to shift from one WAN to another in case of failure.

IPsec

IPsec is fully compatible with multi-WAN. A static route is automatically added for the remote tunnel peer address pointing to the specified WAN gateway to ensure the firewall sends traffic out the correct path when it initiates a connection. For mobile connections, the client always initiates the connection, and the reply traffic is correctly routed by the state table.

An IPsec tunnel may also be set using a gateway group as its interface for failover. This is discussed further in *Multi-WAN Environments*.

OpenVPN

OpenVPN multi-WAN capabilities are described in *OpenVPN and Multi-WAN*. Like IPsec, it can use any WAN or a gateway group.

CARP and multi-WAN

CARP is multi-WAN capable so long as all WAN interfaces use static IP addresses and there are at least three public IP addresses available per WAN. This is covered in *High Availability Configuration Example with Multi-WAN*.

13.3.4 IPv6 and Multi-WAN

IPv6 is also capable of performing in a multi-WAN capacity, but will usually require Network Prefix Translation (NPt) on one or more WANs. This is covered in more detail in *Configuring Multi-WAN for IPv6*.

13.4 Summary of Multi-WAN Requirements

Before covering the bulk of multi-WAN specifics, here is a short summary of the requirements to make a fully implemented multi-WAN setup:

- Create a gateway group under System > Routing on the Groups tab
- Set a failover gateway group for the default gateway as described in *Managing the Default Gateway* (Technically optional but a best practice)
- Configure the DNS Resolver or Forwarder for Multi-WAN, starting by:
  - Use a failover gateway group for the default gaetway (DNS Resolver in default resolver mode)

  - Set at least one unique DNS server for each WAN gateway under System > General Setup with a gateway set (DNS Resolver in forwarding mode or DNS Forwarder)

  - Use the gateway group on LAN firewall rules

13.5 Load Balancing and Failover with Gateway Groups

A Gateway Group is necessary to setup a Load Balancing or Failover configuration. The group itself does not cause any action to be taken, but when the group is used later, such as in policy routing firewall rules, it defines how the items utilizing the group will behave.

The same gateway may be included in multiple groups so that several different scenarios can be configured at the same time. For example, some traffic can be load balanced, and other traffic can use failover, and the same WAN can be used in both capacities by using different gateway groups.

A common example setup for a two WAN firewall contains three groups:

- LoadBalance: Gateways for WAN1 and WAN2 both on Tier 1
- PreferWAN1: Gateway for WAN1 on Tier 1, and WAN2 on Tier 2
- PreferWAN2: Gateway for WAN1 on Tier 2, and WAN2 on Tier 1

No matter which strategy is chosen, the best practice is to have at least one failover group and to set that failover group to be used as the default gateway on the firewall. This ensures that the firewall always has a viable default gateway, and using a gateway group ensures that the correct gateways are used for this function and in the intended order. See *Managing the Default Gateway* for details.

13.5.1 Configuring a Gateway Group for Load Balancing or Failover

To create a gateway group for Load Balancing or Failover:

- Navigate to System > Routing, Groups tab
- Click + Add to create a new gateway group
- Fill in the options on the page as described in *Gateway Group Options*
- Click Save

Load Balancing

Any two gateways on the same tier are load balanced. For example, if *Gateway A*, *Gateway B*, and *Gateway C* are all Tier 1, connections would be balanced between them. Gateways that are load balanced will automatically failover between each other. When a gateway fails it is removed from the group, so in this case if any one of A, B, or C went down, the firewall would load balance between the remaining online gateways.

18.5. Load Balancing and Failover with Gateway Groups

Weighted Balancing

If two WANs need to be balanced in a weighted fashion due to differing amounts of bandwidth between them, that can be accommodated by adjusting the Weight parameter on the gateway as described in *Unequal Cost Load Balancing* and *Advanced Gateway Settings*.

Failover

Gateways on a lower number tier are preferred by the firewall, and if they are down then gateways of a higher numbered tier are used. For example, if *Gateway A* is on Tier 1, *Gateway B* is on Tier 2, and *Gateway C* is on Tier 3, then *Gateway A* would be used first. If *Gateway A* goes down, then *Gateway B* would be used. If both *Gateway A* and *Gateway B* are down, then *Gateway C* would be used.

Complex/Combined Scenarios

By extending the concepts above for Load Balancing and Failover, complicated scenarios are possible that combine both load balancing and failover. For example, if *Gateway A* is on Tier 1, and *Gateway B* and *Gateway C* are on Tier 2, then *Gateway D* on Tier 3, the following behavior occurs: *Gateway A* is preferred on its own. If *Gateway A* is down, then traffic would be load balanced between *Gateway B* and *Gateway C*. Should either *Gateway B* or *Gateway C* go down, the remaining online gateway in that tier would still be used. If *Gateway A*, *Gateway B*, and *Gateway C* are all down, traffic would fail over to *Gateway D*.

Any other combination of the above can be used, so long as it can be arranged within the limit of 5 tiers.

13.5.2 Problems with Load Balancing

Some websites store session information including the client IP address, and if a subsequent connection to that site is routed out a different WAN interface using a different public IP address, the website will not function properly. This is becoming more common with banks and other security-minded sites. One method of working around this issue is to create a failover group and direct traffic destined to these sites to the failover group rather than a load balancing group. Alternately, perform failover for all HTTPS traffic.

The sticky connections feature of pf is intended to resolve this problem, but it has historically been problematic. It is safe to use, and should alleviate this, but there is also a downside to using the sticky option. When using sticky connections, an association is held between the *client IP address* and a given *gateway*, it is not based off of the destination. When the sticky connections option is enabled, any given client would not load balance its connections between multiple WANs, but it would be associated with whichever gateway it happened to use for its first connection. Once all of the client states have expired, the client may exit a different WAN for its next connection, resulting in a new gateway pairing. As such, it works best in environments with many clients where one client utilizing a single WAN does not have a large impact.

13.6 Interface and DNS Configuration

The first two items to configure for Multi-WAN are Interfaces and DNS.

13.6.1 Interface Configuration

Setup the primary WAN as previously described in *Setup Wizard*. Then for the additional WAN interfaces, perform the following tasks:

- Assign the interfaces if they do not yet exist
- Visit the Interfaces menu entry for each additional WAN (e.g. Interfaces > OPT1)
- Enable the interface
- Enter a suitable name, such as WAN2
- Select the desired type of IP address configuration depending on the Internet connection type.
- Enter the remaining details for the type of WAN. For example, on static IP connections, fill in the IP address, subnet mask, and add or select a gateway.

13.6.2 DNS Server Configuration

If the DNS Resolver will be used in forwarding mode or if the DNS Forwarder is in use, the firewall must be configured with DNS servers from each WAN connection to ensure it is always able to resolve DNS. This is especially important if the internal network uses the firewall for DNS resolution.

If the firewall configuration only includes DNS servers from a single WAN, an outage of that WAN connection will result in a complete Internet outage regardless of policy routing configuration since DNS will no longer function.

13.6.3 DNS Resolver Configuration

The DNS Resolver can work with Multi-WAN but the exact configuration depends on the desired behavior and current settings.

If the DNS Resolver must work in its default resolver mode, such as for environments which require DNSSEC, then forwarding mode cannot be enabled. This can still function with Multi-WAN but requires using failover for the default gateway. See *Managing the Default Gateway*.

If the DNS Resolver can use forwarding mode, then the following procedure may be performed instead:

- Set at least one DNS server per WAN under System > General Setup, as described in the next section
- Check Enable Forwarding Mode under Services > DNS Resolver
- Uncheck Enable DNSSEC Support

DNS Servers and Static Routes

When using the DNS Forwarder or the DNS Resolver in forwarding mode, the firewall uses its routing table to reach the configured DNS servers. This means without any static routes configured, it will only use the primary WAN connection to reach DNS servers. Gateways must be selected for each DNS server defined on the firewall to use the correct WAN interface to reach that DNS server. DNS servers that come from dynamic gateways are automatically routed back out the proper path. At least one gateway from each WAN should be selected where possible.

To configure the DNS server gateways:

- Navigate to System > General Setup

- Define at least one *unique* DNS server for each WAN
- For each DNS server, select an appropriate gateway so it uses a specific WAN interface

18.6. Interface and DNS Configuration

Note: The same DNS server cannot be entered more than once. Each entry must be unique.

Selecting gateways for DNS servers is required for several reasons. One, most ISPs prohibit recursive queries from hosts outside their network, hence the firewall must use the correct WAN interface when accessing DNS servers for a specific ISP. Secondly, if the primary WAN fails and the firewall does not have a gateway chosen for one of the other DNS servers, the firewall will lose all DNS resolution ability from the firewall itself. Access to DNS is lost in that situation because all DNS servers will be unreachable when the default gateway is unreachable. If the firewall is used as a DNS server for the local network, this will result in a complete failure of DNS.

When using the DNS Resolver with forwarding mode disabled, the unbound daemon speaks directly to the root DNS servers and other authoritative DNS servers, which makes using such static routes and gateway assignments impossible. In that case, configure failover for the default gateway (*Managing the Default Gateway*) so that the DNS Resolver can maintain outbound connectivity.

13.6.4 Scaling to Large Numbers of WAN Interfaces

There are numerous users of AZTCO-FW® software deploying 6-12 Internet connections on a single installation. One user has 10 DSL lines because in his country it is significantly cheaper to get ten 256 Kb connections than it is one 2.5 Mb connection. That customer load balances a large number of internal machines out 10 different connections. For more information on this scale of deployment, see *Multi-WAN on a Stick* later in this chapter.

13.7 Multi-WAN and NAT

The default NAT rules generated by AZTCO-FW® software will translate any traffic leaving a WAN-type interface to the IP address of that interface. In a default two interface LAN and WAN configuration, AZTCO-FW will NAT all traffic from the LAN subnet leaving the WAN interface to the WAN IP address. Adding more WAN-type interfaces extends this to NAT any traffic leaving a WAN-type interface to that interface IP address. This is all handled automatically unless Manual Outbound NAT is enabled.

Warning: NAT does not influence the path taken by connections, only how addresses on packets traversing an interface are translated by the firewall.

Policy routing firewall rules direct connections to specific WAN interfaces, and the Outbound and 1:1 NAT rules specify how the addresses on packets for those connections will be translated by the firewall as it leaves that WAN.

13.7.1 Multi-WAN and Manual Outbound NAT

If Manual Outbound NAT must be used with multi-WAN, ensure manual outbound NAT rules are present on all WAN-type interfaces.

13.7.2 Multi-WAN and Port Forwarding

Each port forward applies to a single WAN interface. A given port can be opened on multiple WAN interfaces by using multiple port forward entries, one per WAN interface. The easiest way to accomplish this is:

- Add a port forward on the first WAN connection as usual
- Click ⊘ to the right of that entry to add another port forward based on the selected one
- Change the Interface to the desired WAN
- Click Save

The reply-to keyword in pf, used on WAN-type interface rules, ensures that when traffic comes in over a specific WAN interface, the return traffic will go back out the way it came into the firewall. So port forwards can be actively used on all WAN interfaces at any time, regardless of any failover configuration that may be present. This is especially useful for mail servers, as an address on a secondary WAN can be used as a backup MX, allowing the site to receive mail even when the primary line is down. This behavior is configurable, for information on this setting, see *Disable Reply-To*.

13.7.3 Multi-WAN and 1:1 NAT

1:1 NAT entries are specific to a single WAN interface and, like outbound NAT, they only control what happens to addresses on packets as they pass through an interface. Internal systems can be configured with a 1:1 NAT entry on each WAN interface, or a 1:1 entry on one or more WAN interfaces and use the default outbound NAT on others. Where 1:1 entries are configured, they always override any other Outbound NAT configuration for that specific interface.

If a local device must always use a 1:1 NAT entry on a specific WAN, then traffic from that device must be forced to use that specific WAN gateway with policy routing firewall rules.

13.8 Policy Routing Configuration

At this point, the firewall is prepared for Multi-WAN but it will not yet be used. Traffic will not properly fail over or be load balanced without policy routing firewall rules in place.

Note: An exception to this is when using a gateway group or automatic failover for the default gateway (*Managing the Default Gateway*). In this case failover could still function without policy routing, but not load balancing.

13.8.1 Configuring Firewall Rules for Policy Routing

Setting a Gateway on a firewall rule will cause traffic matching the rule to use the chosen gateway or group, following the configured behavior of the group.

The easiest way to configure a firewall for policy routing is to edit the existing default pass rule for the LAN and select the gateway group there. With that set, any traffic matching the default pass rule on the LAN will use the chosen gateway or group. To make that edit:

- Navigate to Firewall > Rules, LAN tab

18.8. Policy Routing Configuration

- Click ⧉ on the row with the default pass rule
- Click ⚙ Display Advanced under Extra Options
- Select the desired gateway group from the Gateway drop-down list
- Click Save
- Click Apply Changes

Only the most basic of deployments will be satisfied with that configuration, most configurations are more complex. Continue reading for more factors that can require additional configuration.

13.8.2 Bypassing Policy Routing

If there are other local interfaces, VPNs, MPLS interfaces, or traffic that must otherwise follow the system routing table, then that traffic must be configured to bypass policy routing. This is simple to do by making a rule to match the traffic in question and then placing that rule above any rules that have a gateway configured, because the first rule to match is the one that is used.

This can be generalized by making an alias for any *RFC1918* traffic which would cover all private networks, and then using that in a rule. The alias contains 192.168.0.0/16, 172.16.0.0/12, and 10.0.0.0/8.

In Figure *Bypass Policy Routing Example Rules*, local and VPN traffic bypasses policy routing, HTTPS traffic prefers WAN2, and all other traffic is load balanced:

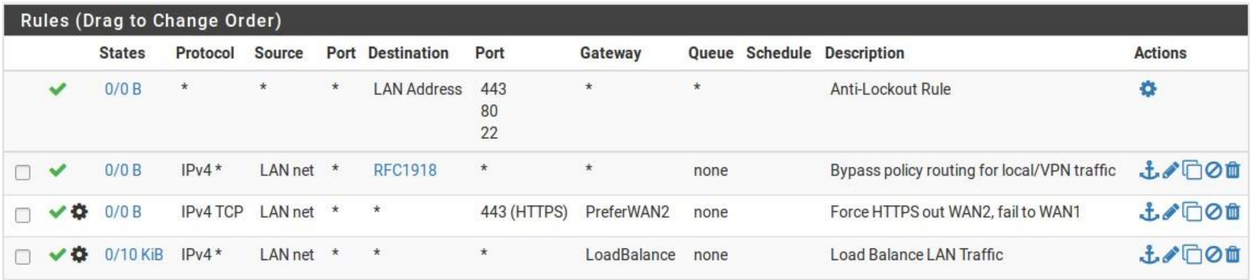

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
	✓	0/0 B	*	*	*	LAN Address	443 80 22	*	*		Anti-Lockout Rule	⚙
□	✓	0/0 B	IPv4 *	LAN net	*	RFC1918	*	*	none		Bypass policy routing for local/VPN traffic	⚓✎⧉⊘🗑
□	✓⚙	0/0 B	IPv4 TCP	LAN net	*	*	443 (HTTPS)	PreferWAN2	none		Force HTTPS out WAN2, fail to WAN1	⚓✎⧉⊘🗑
□	✓⚙	0/10 KiB	IPv4 *	LAN net	*	*	*	LoadBalance	none		Load Balance LAN Traffic	⚓✎⧉⊘🗑

Fig. 1: Bypass Policy Routing Example Rules

13.8.3 Mixing Failover and Load Balancing

As shown in Figure *Bypass Policy Routing Example Rules*, failover and load balancing can be used at the same time by carefully ordering the rules on an interface. Rules are processed from the top down and the first match wins. By placing more specific rules near the top of the list, and the general "match all" style rules at the bottom, any number of different combinations are possible with rules using different gateways or groups.

13.8.4 Enforcing Gateway Use

There are situations where traffic should only ever use one gateway and never load balance or failover. In this example, a device must only exit via a specific WAN and lose all connectivity when that WAN fails.

First, set the Gateway on a firewall rule matching traffic from this device to a specific WAN Gateway. If that gateway is down, the rule will act as if the gateway was not set at all, so it needs to be taken a couple steps further.

Add a rule immediately below the rule matching the traffic, but set to reject or block instead. This rule must not have a gateway set.

Next, configure the firewall to omit rules for gateways that are down (*Gateway Monitoring*):

- Navigate to System > Advanced on the Miscellaneous tab
- Check Do not create rules when gateway is down
- Click Save

With that option enabled, the first rule will be omitted entirely, falling through to the next matching rule. This way, when the first rule is omitted automatically, traffic will be stopped by the block rule.

13.9 Verifying Functionality

Once multi-WAN has been configured, best practice is then to test its functionality to verify it functions as expected. The following sections describe how to test each portion of a multi-WAN configuration.

13.9.1 Testing Failover

Testing Multi-WAN in a controlled manner immediately after configuration is a key step in the process. Do not make the mistake of waiting until an Internet connection fails naturally for the first test, only to discover problems when they are much more difficult and stressful to fix.

First, navigate to Status > Gateways and ensure all WAN gateways are show as Online under Status, as well as on the Gateway Groups tab. If they do not, verify that a proper monitor IP address is used as discussed in *Gateway Settings*.

Creating a WAN Failure

There are a number of ways to simulate a WAN failure depending on the type of Internet connection being used. For any type, first try unplugging the target WAN interface Ethernet cable from the firewall.

For cable and DSL connections, try powering off the modem/CPE, and in a separate test, unplug the coax or phone line from the modem. For fiber, wireless, and other types of connections with a router outside of AZTCO-FW® software, try unplugging the Internet connection from the router, and also turning off the router itself.

All of the described testing scenarios will likely end with the same result. However, there are some circumstances where trying all these things individually will find a fault that would not have otherwise been noticed until an actual failure. One of the most common is unknowingly using a monitor IP address assigned to the DSL or cable modem. Hence when the coax or phone line is disconnected, simulating a provider failure rather than an Ethernet or modem failure, the monitor ping still succeeds since it is pinging the modem. From what the firewall was told to monitor, the connection is still up, so it will not fail over even if the upstream connection is actually down. There are other types of failure that can similarly only be detected by testing all the individual possibilities for failure. The monitor IP address can be edited on the gateway entry as covered in *Gateway Settings*.

18.9. Verifying Functionality

Verifying Interface Status

After creating a WAN failure, refresh Status > Gateways to check the current status. As the gateway monitoring daemon notices the loss, the loss will eventually move past the configured alarm thresholds and it will mark the gateway as down.

13.9.2 Verifying Load Balancing Functionality

This section describes how to verify the functionality of a load balancing configuration.

Testing load balancing with traceroute

The traceroute utility (or tracert in Windows) shows the network path taken to a given destination. See *Using traceroute* for details on using traceroute. With load balancing, running a traceroute from a client system behind the firewall should show a different path being taken for each attempt. Due to the way traceroute functions, wait at least one minute after stopping a traceroute before starting another test so that the states will expire, or try different destinations on each attempt.

18.9. Verifying Functionality

Using Traffic Graphs

The real time traffic graphs under Status > Traffic Graph and on the Traffic Graphs dashboard widget are useful for showing the real time throughput on WAN interfaces. Only one graph at a time can be shown per browser window when using Status > Traffic Graph, but additional windows or tabs can be opened in the browser to see all WAN interfaces simultaneously. The traffic graphs widget for the Dashboard enables the simultaneous display of multiple traffic graphs on a single page to simplify this process. If load balancing is working correctly, activity will be observed on all WAN interfaces.

The RRD traffic graphs under Status > Monitoring are useful for longer-term and historical evaluation of WAN utilization.

Note: Bandwidth usage may not be exactly equally distributed, since connections are directed on a round robin basis without regard for bandwidth usage.

13.10 Multi-WAN on a Stick

In the router world, Cisco and others refer to a VLAN router as a “router on a stick” since it can be a functioning router with only one physical network connection. AZTCO-FW® software can be configured in this manner as well, using VLANs and a managed switch capable of 802.1q trunking. Most of the deployments running more than 5 WANs use this methodology to limit the number of physical interfaces required on the firewall. In such a deployment, the WAN interfaces all reside on one physical interface on the firewall, with the internal network(s) on additional physical interfaces. Figure *Multi-WAN on a stick* illustrates this type of deployment.

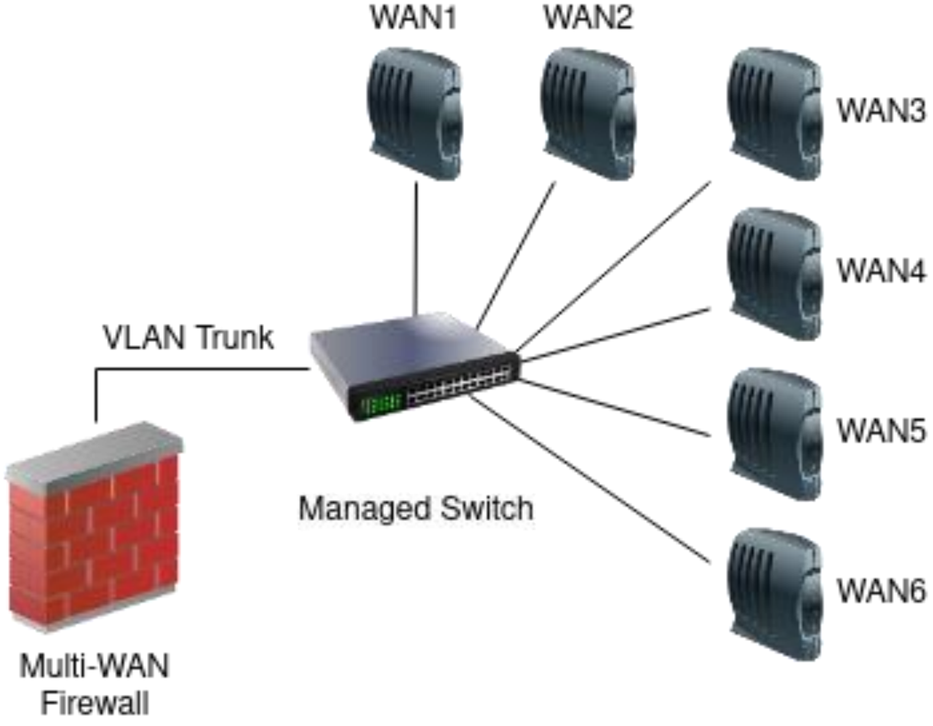

Fig. 2: Multi-WAN on a stick

13.11 Multi-Link PPPoE (MLPPP)

Multi-Link PPPoE (MLPPP) is a unique WAN option that bonds together multiple PPPoE lines from the same ISP to form one larger virtual circuit. This means a firewall can get the true aggregate bandwidth of all circuits in the bundle. For example, if a firewall has three 10 Mbit/s DSL lines in a bundle, it could potentially get 30Mbit/s from a single connection.

13.11.1 Requirements

The largest hurdle for MLPPP is that the ISP must support it on the circuits connected to the firewall. Few ISPs are willing to support MLPPP, so if an ISP is available that does, it would be worth taking advantage of that fact. Additionally, each line must be on a separate interface connected to the firewall running AZTCO-FW® software.

13.11.2 Setup

Setup for MLPPP is simple:

- Configure a WAN for a single line with the correct credentials
- Navigate to Interfaces > Assign, PPPs tab
- Click ⧉ to edit the entry for the PPPoE WAN
- Ctrl-click to select the other physical interfaces that belong to the same MLPPP bundle
- Click Save

The firewall will then attempt to bond the lines using MLPPP.

13.11.3 Caveats

One downside to using MLPPP is that troubleshooting is much more difficult. Statistics and status are not available for the individual lines. To determine the status, read through the PPP log, as there is not yet a way to query the lines separately. In some cases it's obvious if a line is down, as there may be a noticeable problem at the modem (out of sync) or that the maximum attainable bandwidth is reduced.

13.12 Using OpenVPN with Multi-WAN

OpenVPN servers can be used with any WAN, or multiple WANs, as can OpenVPN clients. This document covers only a remote access OpenVPN server, but a similar process could be applied for site to site VPNs.

18.12.1 OpenVPN Configuration

First, get OpenVPN working as desired on the primary WAN interface. Once it is properly functioning, make a backup.

13.12.2 Bind to Localhost and Setup Port Forwards

The OpenVPN configuration needs to be adjusted so it can be reached from either WAN. The simplest way to do this is by changing the Interface on the VPN connection to be *Localhost*, and then adding a port forward on each WAN to redirect the OpenVPN port to *Localhost* (127.0.0.1).

For example: If there are two WANs and the OpenVPN server is running on port *1194*, set the Interface to *Localhost*, then add two port forwards:

- WAN1 - *UDP*, Source any, Destination *WAN1 Address* port 1194, redirect target 127.0.0.1 port 1194
- WAN2 - *UDP*, Source any, Destination *WAN2 Address* port 1194, redirect target 127.0.0.1 port 1194

13.12.3 Configure Clients

Clients may be configured to use the second WAN by adding a second *remote* statement to their configuration, such as:

```
remote x.x.x.x 1194 udp
```

Where x.x.x.x is the second WAN IP address or host name.

This process can be automated by using the OpenVPN Client Export package. When exporting a client, in Host Name Resolution choose one of:

Automagic Multi-WAN IPs (port forward targets) Adds a remote statement for each port forward found targeting the interface binding and port used by this VPN, uses the IP address of each WAN as-is.

Automagic Multi-WAN DDNS Hostnames (port forward targets) Like above, but uses the first located Dynamic DNS hostname for a given WAN. If the WAN is a private IP address, this may be the better choice.

13.12.4 More than two WAN connections

The same steps can be repeated to add more WAN connections. Add a port forward to any additional WAN. Clients will need an updated configuration file if another WAN is added later.

See also:

- *Troubleshooting Multi-WAN*
- *Configuring Multi-WAN for IPv6*

The multiple WAN (multi-WAN) capabilities in AZTCO-FW® software allow a firewall to utilize multiple Internet connections to achieve more reliable connectivity and greater throughput capacity.

> Warning: Before proceeding with a multi-WAN configuration, the firewall must have a functional two interface (LAN and WAN) configuration.

18.12. Using OpenVPN with Multi-WAN

AZTCO-FW software is capable of handling numerous WAN interfaces, with multiple deployments using over 10 WANs in production.

All WAN-type interfaces are treated identically in the GUI. Anything that can be done with the primary WAN can also be done with an additional OPT WAN interface. There are no significant differences between the primary WAN and additional WANs.

This section starts by covering items to consider when implementing *any* multi-WAN solution, then covers multi-WAN configuration with AZTCO-FW software.

13.13 Choosing Internet Connectivity

The ideal choice of Internet connectivity will depend largely upon the options available at a given location, but there are some additional factors to take into consideration.

13.13.1 Cable Paths

Speaking from the experience of those who have seen first hand the effects of multiple cable-seeking backhoes, as well as nefarious copper thieves, it is highly desirable to obtain connectivity choices for a multi-WAN deployment which utilize disparate cabling paths. In many locations, DSL connections as well as any others utilizing copper pairs are carried on a single cable subject to the same cable cut, and others from the same telco such as fiber circuits may run along the same poles or conduits.

If one connection comes in over copper pair (DSL), choose a secondary connection utilizing a different type and path of cabling. Cable connections are typically the most widely available option not subject to the same outage as copper services. Other options include fiber service or fixed wireless coming in on a different path from copper services.

Two connections of the same type cannot be relied upon to provide redundancy in most cases. An ISP outage or cable cut will commonly take down all connections of the same type. Some users use multiple DSL lines or multiple cable modems, though the only redundancy that typically offers is isolating a site from modem or other CPE (Customer Premise Equipment) failure. Consider multiple connections from the same provider as a solution only for additional bandwidth, as the redundancy such a deployment offers is minimal.

13.13.2 Paths to the Internet

Another consideration when selecting Internet connectivity for a site is the path from the connection itself to the Internet. For redundancy purposes, multiple Internet connections from the same provider, especially of the same type, should not be relied upon as they could all fail concurrently.

With larger providers, two different types of connections such as a Fiber line and DSL will usually traverse significantly different networks until reaching core parts of the network. These core network components are generally designed with high redundancy and any problems are addressed quickly since they have widespread effects. Hence such connectivity is isolated from most ISP issues, but since they commonly utilize the same cable path, it still leaves a site vulnerable to extended outages from cable cuts.

13.13.3 Better Redundancy, More Bandwidth, Less Money

In the past, high-grade telco services such as DS1 or DS3 circuits were the choice for environments with high availability requirements. Generally the Service Level Agreements (SLA) offered on DS1 and DS3 connections were better than other types of connectivity, and those circuits were generally seen as more reliable. End-users have largely left such circuits behind, however, because they are too slow or too costly by today's standards. With the multi-WAN capabilities on AZTCO-FW, a site can have more bandwidth and better redundancy for less money in many cases. Fiber services are rapidly becoming more widespread, shaking up this concept by providing extremely large amounts of bandwidth for relatively low cost, though such services may still have a less-than-desirable SLA for outage response.

Most organizations requiring high availability Internet connections do not want to rely upon DSL, cable or other "lesser class" broadband Internet connections. While they're usually significantly faster and cheaper, the lesser SLA is enough to make many companies think twice. In areas where multiple lower cost broadband options are available, such as fiber and cable, the combination of AZTCO-FW software and two low cost Internet connections provides more bandwidth and better redundancy at a lower cost. The chance of two different broadband connections going down simultaneously is significantly less than the chance of any single service outage. Adding a backup Cable or DSL line to supplement a much faster fiber line ensures connectivity will continue when an outage occurs on the fiber line, even if it is a rare occurrence.

CHAPTER

FOURTEEN

VIRTUAL PRIVATE NETWORKS

14.1 Choosing a VPN solution

Each VPN solution has pros and cons. This section will cover the primary considerations in choosing a VPN solution, providing the information necessary to choose the best solution for a given environment.

14.1.1 Interoperability

To interoperate with a firewall or router product from another vendor, IPsec is usually the best choice since it is included with nearly every VPN-capable device. It also prevents being locked into any particular firewall or VPN solution. For interoperable site-to-site connectivity, IPsec is usually the only choice. OpenVPN is interoperable with a few other packaged firewall/VPN solutions, but not many. Interoperability in this sense isn't applicable with other VPN types since they are not intended for site-to-site applications.

14.1.2 Authentication considerations

In current versions of AZTCO-FW® software, all VPN types support user authentication. IPsec and OpenVPN can also work with shared keys or certificates. OpenVPN is a bit more flexible in this regard because it can work with only certificates, only shared keys, only user authentication, or a combination of these. Using OpenVPN with certificates, TLS authentication, and User Authentication is the most secure method. OpenVPN certificates can also be password protected, in which case a compromised certificate alone isn't adequate for connecting to a VPN if it is set to only use certificates. The lack of additional authentication can be a security risk in that a lost, stolen, or compromised system containing a key or certificate means whoever has access to the device can connect to a VPN until that loss is discovered and the certificate revoked.

While not ideal, a lack of username and password authentication on a VPN isn't as great a risk as it may seem. A compromised system can easily have a key logger installed to capture the username and password information and easily defeat that protection. In the case of lost or stolen systems containing keys, if the hard drive isn't encrypted, the keys can be used to connect. However adding password authentication isn't of great help there either, as usually the same username and password will be used to log into the computer, and most passwords are crackable within minutes using modern hardware when an attacker has access to an unencrypted drive. Password security is also frequently compromised by users with notes on their laptop or in their laptop case with their password written down. As with any security implementation, the more layers utilized, the better, but it's always a good idea to keep these layers in perspective.

14.1.3 Ease of configuration

None of the available VPN options are extremely difficult to configure, but there are differences between the options:

- IPsec has numerous configuration options and can be difficult for the uninitiated.
- OpenVPN requires the use of certificates for remote access in most environments, which comes with its own learning curve and can be a bit arduous to manage. AZTCO-FW includes a wizard to handle the most common OpenVPN remote access configurations and the OpenVPN client export packages eases the process of getting the clients up and running.

IPsec and OpenVPN are preferable options in many scenarios for other reasons discussed throughout this chapter.

14.1.4 Multi-WAN capable

If users require the ability to connect to multiple WANs, both IPsec and OpenVPN are capable of handling such configurations.

14.1.5 Client availability

VPN Client software is a program that handles connecting to the VPN and handling any other related tasks like authentication, encrypting, routing, etc. For remote access VPNs, the availability of VPN client software is a primary consideration. All options are cross platform compatible with many different operating systems but some require installing third-party clients. IPsec in EAP-MSCHAPv2 mode, IPsec in EAP-TLS mode, and IPsec in Xauth mode are the only options with client support built into some popular desktop and mobile operating systems. Other operating systems vary and may include more or less IPsec modes or may even include OpenVPN, as is the case with many Linux distributions. If using built-in clients is a must, consult the operating system documentation for all required client platforms to see if a common option is available and then check AZTCO-FW to see if that mode is possible.

In some cases multiple remote access VPNs may be required to accommodate all clients. For example, IPsec could be used for some and OpenVPN for others. Some organizations prefer to keep things consistent, so there is a trade-off to be made but for the sake of compatibility it may be worth offering multiple options.

IPsec

IPsec clients are available for Windows, Mac OS X, BSD, Linux, and others. Though the native clients may only support certain specific modes and configurations. General-use IPsec clients are not included in the OS except for some Linux and BSD distributions. A good free option for Windows is the Shrew Soft client. Mac OS X includes both IKEv2 and Cisco (xauth) IPsec support. There are free and commercial options available with a user-friendly GUI.

OSX 10.11, along with Windows 7 and later include support for IPsec in specific modes using IKEv2: EAP-TLS and EAP-MSCHAPv2. Both options are supported by AZTCO-FW and are covered in *IPsec*.

The Cisco-style IPsec client included with OS X and iOS devices is fully compatible with AZTCO-FW IPsec using xauth. Configuration for the iOS client is covered in *Configuring IPsec IKEv2 Remote Access VPN Clients on iOS*.

Many Android phones also include a compatible IPsec client, which is discussed in *Configuring IPsec IKEv2 Remote Access VPN Clients on Android*.

14.1. Choosing a VPN solution

OpenVPN

OpenVPN has clients available for Windows, Mac OS X, all the BSDs, Linux, Solaris, and Windows Mobile, but the client does not come pre-installed in any of these operating systems.

Android 4.x and later devices can use a freely available OpenVPN client that works well and doesn't require rooting the device. That client is covered in *Installing the OpenVPN Client on Android*. Older versions of Android may also be able to use OpenVPN via an alternate client. There are other options available if the device is rooted, but that is beyond the scope of this documentation. iOS also has a native OpenVPN client. For more information, see *Installing the OpenVPN Client on iOS*.

14.1.6 Firewall friendliness

VPN protocols can cause difficulties for many firewalls and NAT devices. This is primarily relevant to remote access connectivity, where users will be behind a myriad of firewalls mostly controlled by third parties with varying configurations and capabilities.

IPsec

IPsec uses both UDP port 500 and the ESP protocol to function. Some firewalls don't handle ESP traffic well where NAT is involved, because the protocol does not have port numbers like TCP and UDP that make it easily trackable by NAT devices. IPsec clients behind NAT may require NAT Traversal to function, which encapsulates the ESP traffic over UDP port 4500.

OpenVPN

OpenVPN is the most firewall-friendly of the VPN options. Since it uses TCP or UDP and is not affected by any common NAT functions such as rewriting of source ports, it is rare to find a firewall which will not work with OpenVPN. The only possible difficulty is if the protocol and port in use is blocked. Some administrators use a common port like UDP 53 (usually DNS), or TCP 80 (usually HTTP) or TCP 443 (usually HTTPS) or to evade most egress filtering.

14.1.7 Cryptographically secure

One of the critical functions of a VPN is to ensure the confidentiality of the data transmitted.

IPsec using pre-shared keys can be broken if a weak key is used. Use a strong key, at least 10 characters in length containing a mix of upper and lowercase letters, numbers and symbols. Use of certificates is preferred, though somewhat more complicated to implement.

OpenVPN encryption is compromised if the PKI or shared keys are disclosed, though the use of multiple factors such as TLS authentication on top of PKI can mitigate some of the danger.

14.1. Choosing a VPN solution

14.1.8 Recap

Table *Features and Characteristics by VPN Type* shows an overview of the considerations provided in this section.

Table 1: Features and Characteristics by VPN Type

VPN Type	Client included in most OSes	Widely interoperable	Multi-WAN	Cryptographically secure	Firewall friendly
IPsec	Varies by mode	Yes	Yes	Yes	No (without NAT-T)
Open-VPN	No	No	Yes	Yes	Yes

14.2 Remote Access Mobile VPN Client Compatibility

AZTCO-FW® software supports a variety of remote access ("mobile") VPN configuration styles to accommodate nearly any potential client. The table below shows which operating systems have compatible clients with some of the most common remote access VPN configurations available on AZTCO-FW software.

Operating System	Protocol						
	OpenVPN	IPsec					
	————— *3PA* 1 2	——————— PSK	——— –RSA	Xauth PSK	Xauth RSA	IKEv2 EAP MSCHAPv2	IKEv2 EAP TLS /RADIUS
Windows XP		3PA 3	3PA 3	*3PA* 3	3PA 3	?	?
Windows Vista/7/8	*3PA* 1 2	3PA 3	3PA 3	*3PA* 3	3PA 3	*Yes (7+)*	*Yes (7+)*
Windows 10	*3PA* 1 2	?	?	?	?	*Yes*	*Yes*
Android <4	*3PA*	Varies	Varies	Bug	Yes	?	?
Android 4+	*3PA* 4	Varies	Varies	Bug	Yes	*3PA* 5	3PA 5
iOS < 9	*3PA* 6	?	?	*Yes*	Yes	?	?
iOS 9+	*3PA* 6	?	?	*Yes*	Yes	*Yes*	*Yes*
OS X < 10.11	*3PA* 2	?	?	*Yes*	Yes	?	?
OS X 10.11+	*3PA* 2	?	?	*Yes*	Yes	*Yes*	*Yes*
SNOM/Yeali	nkYes	No	No	No	No	No	No

Table: Mobile/Remote Access VPN Client Availability

- Yes = OS Native Client Available
- 3PA = Third Party Client Required
- Bug = Known problem configuration, follow the link for more details
- Varies = Varies by device model and vendor options

Unless otherwise stated, UNIX clients (*BSD, Linux, etc) can support any style with manual configurations but the availability of GUI configuration tools varies by distribution.

19.2. Remote Access Mobile VPN Client Compatibility

14.3 VPNs and Firewall Rules

VPNs and firewall rules are handled somewhat inconsistently in AZTCO-FW® software. This section describes how firewall rules are handled for each of the individual VPN options. For the automatically added rules discussed here, the addition of those rules may be disabled by checking Disable all auto-added VPN rules under System > Advanced on the Firewall/NAT tab.

14.3.1 IPsec

IPsec traffic coming in to the specified WAN interface is automatically allowed as described in *IPsec*. Traffic encapsulated within an active IPsec connection is controlled via user-defined rules on the IPsec tab under Firewall > Rules.

19.3.2 OpenVPN

OpenVPN does not automatically add rules to WAN interfaces. The OpenVPN remote access VPN Wizard offers to optionally create rules to pass WAN traffic and traffic on the OpenVPN interface. Traffic encapsulated within an active OpenVPN connection is controlled via user-defined rules on the OpenVPN tab under Firewall > Rules. OpenVPN interfaces may also be assigned similar to other interfaces on AZTCO-FW. In such cases the OpenVPN tab firewall rules still apply, but there is a separate tab specific to the assigned VPN instance that controls traffic only for that one VPN.

14.4 VPNs and IPv6

There are some special considerations for VPNs when using them in combination with IPv6. The two main items of concern are:

- Whether or not a certain VPN type supports IPv6
- Making sure the firewall rules don’t allow unencrypted traffic in that should be coming over a VPN.

14.4.1 IPv6 VPN Support

Support for IPv6 varies from type to type and in client support. Be sure to check with the vendor of the other device in order to make sure a non-AZTCO-FW firewall or client supports IPv6 VPNs.

IPsec

AZTCO-FW® software supports IPsec using IKEv1 over IPv6 with one quirk: If a tunnel uses an IPv6 peer address, the tunnel can only carry IPv6 phase 2 networks, and the same for IPv4. Traffic cannot be mixed between address families with IKEv1. See *IPsec and IPv6*.

When an IPsec tunnel is set for IKEv2, it can include both IPv4 and IPv6 Phase 2 definitions concurrently.

OpenVPN

OpenVPN fully supports IPv6 for site-to-site and mobile clients, and tunnels can carry both IPv4 and IPv6 traffic concurrently. See *OpenVPN and IPv6*.

14.4.2 IPv6 VPN and Firewall Rules

As mentioned briefly in *Firewall and VPN Concerns*, special care must be taken when routing IPv6 traffic across a VPN and using publicly routable subnets. The same advice also applies to IPv4 but it's much less common to have clients on both sides of an IPv4 VPN using publicly routable addresses.

The main issue is that because it's possible to route all the way from one LAN to the other LAN across the Internet, then traffic could be flowing unencrypted between the two networks if the VPN is down (or not present at all!). This is far from ideal because although connectivity is available, if any traffic were intercepted in between the two networks and that traffic was using an unencrypted protocol like HTTP, then it could compromise the network.

One way to prevent this is to not allow traffic from the remote IPv6 LAN in on the opposing side's WAN rules. Only allow traffic from the remote side's subnet on the firewall rules for whichever VPN type is being used to protect the traffic. An explicit block rule could also be added to the top of the WAN rules to ensure that this traffic cannot enter from the WAN directly. A better method is to use a floating rule to reject outbound traffic on WAN destined for VPN hosts/remote local networks. This way the insecure traffic never leaves the premises. With the rule set to log, the "leakage" would be obvious to someone monitoring the logs as it would be shown blocked outbound on WAN.

Another less obvious consequence of having dual stack connectivity between networks is that differences in DNS can cause unintended routing to take place. Suppose IPv4 VPN connectivity exists between two sites, but there is no IPv6 VPN, only standard IPv6 connectivity at both locations. If a local host is set to prefer IPv6 and it receives a AAAA DNS response with the IPv6 IP address for a remote resource, it would attempt to connect over IPv6 first rather than using the VPN. In cases such as this, care would be needed to make sure that DNS does not contain conflicting records or that floating rules are added to prevent this IPv6 traffic from leaking out WAN. A more in-depth article on these kinds of traffic leakage can be found in the IETF draft named draft-gont-opsec-vpn-leakages-00.

14.5 VPN Scaling

The advice on this page is intended to help firewall administrators handle increased VPN volume, both in terms of throughput and number of connected users.

Warning: The advice on this page is relayed from experience and from community members. The advice on this page may not apply to all environments or use cases, and has not been definitively proven to help, but is offered in case others find it useful.

14.5.1 General Advice

No AZTCO-FW Limits

AZTCO-FW does not place any artificial limits on VPN connections. Any limitations encountered are due to settings, the hardware/environment, or the underlying technology.

IPsec is Faster

IPsec is faster than OpenVPN, so if both client and server support IPsec, use IPsec.

Use External Authentication

For user-based authentication, the most efficient method of user management for large numbers of accounts is an external authentication source, such as a RADIUS server, LDAP server, Active Directory (Via LDAP or RADIUS/NPS), etc.

Check Logs

If additional users are unable to connect, look in the logs on both the client and server side for specific error messages before seeking support.

Use Hardware Acceleration

Using a cryptographic accelerator such as a CPU with AES-NI will help greatly with throughput and crypto-related tasks. Enable the AES-NI and BSD cryptodev modules under System > Advanced on the Miscellaneous tab.

Use AES-GCM

Using efficient encryption like AEAD ciphers, which combine encryption and authentication, will increase security and performance. Both IPsec and OpenVPN can use AES-GCM, which is an AEAD cipher. Client support may vary by platform.

Use Accelerated Ciphers

As above, using AES-NI and AES-GCM is the best possible combination at this time. That said, certain hardware may accelerate other ciphers so that alternate choices are faster or more efficient. For hardware sold by AZTCO-FW, see the AZTCO-FW Appliances page for performance data and recommendations.

Disable Performance-Limiting Mitigation Settings

While we do not have any data on if or by how much they may impact VPNs, CPU vulnerability mitigation methods such as Kernel PTI and MDS mode can potentially degrade total performance. The potential for exploitation is minimal since arbitrary code cannot be run on the firewall except by users which already have the equivalent of administratorlevel access. To ensure this risk stays low, only allow trusted administrators to access the firewall GUI and shell (SSH or console). The settings to enable/disable these features are under System > Advanced on the Miscellaneous tab.

Check Tunnel Network/Virtual Address Pool Sizes

Both IPsec and OpenVPN can assign addresses to clients out of a pool for remote access/mobile VPNs. The sizing of this pool limits how many clients can connect. For example, the maximum number of users in a /24 pool is 252, but other settings may reduce that value. See the sections below for more specific advice.

Use "Secure Enough" Settings

While we do not recommend deliberately using weak configurations, in some cases trade-offs are made for security between two secure ciphers or settings where one may offer *even better* security, but the lower of the two is still secure. In these cases, using the "Secure Enough" option can provide efficiency vs increased security. So long as the decision is informed, there may be some performance gained without compromising security in an unacceptable way. For example, with AES-GCM a key length of 128 bits is still considered secure. A 256 bit key is more secure, but the 256 bit key could put more of a burden on the hardware.

Consider Split Tunneling

Configurations which send all client data over the VPN, including Internet-bound traffic, will consume more resources than those which only send traffic for specific subnets. There are plenty of valid reasons to use either kind of configuration, however, when resources are stretched thin, easing the traffic burden on the VPN may justify switching to split tunneling rather than tunneling everything. Depending on the type of VPN and client, this may require adjustments on the server, the client, or both. See the sections below for specific recommendations.

Use Multiple Firewalls

In some instances, the burden may be too great for any single AZTCO-FW firewall to handle. In cases like this, multiple firewalls can be used to handle the required number of clients or throughput, at a cost of greatly increased complexity. There is no way to automatically balance between nodes in this manner, but such a configuration could be manually managed. This would also likely require the capability to have multiple external addresses on the WAN so each firewall can work in parallel, and also increases the complexity of routing on the internal side.

14.5.2 Scaling IPsec

IPsec is well-suited to high throughput by default, especially given the advice above, but there are additional IPsecspecific tweaks which may help.

Optimal Encryption Settings

- Use AES-NI capable hardware.
- In Phase 1 (IKE) settings, use:
  - *AES128-GCM* with *128 bit* key length for the Algorithm
  - *AES-XCBC* for the hash, which in this case is effectively a Pseudo-Random Function (PRF).
- In Phase 2 (Child SA) settings, use:
  - *AES128-GCM* with *128 bit* key length for the Algorithm
  - Do not select any Hash Algorithms. A hash algorithm is unnecessary for AES-GCM as it already includes authentication.

Enable Multiple Phase 1 and Phase 2 Proposals

Multiple Phase 1 and Phase 2 encryption proposals may be configured in the GUI. Enabling multiple combinations of settings will allow peers to choose the most optimal settings which both sides support.

Enable Asynchronous Cryptography

IPsec cryptography jobs can be dispatched multi-threaded to run in parallel and increase performance. However, not all platforms and configurations fully support this function. To enable this capability, check Asynchronous Cryptography under VPN > IPsec on the Advanced tab.

Split Tunneling

As mentioned above, split tunneling would only send traffic for specific subnets across the VPN rather than sending all traffic. On IPsec, this can be done in some cases by listing the specific networks in Phase 2 entries for the Mobile IPsec P1 rather than 0.0.0.0/0. On the mobile clients tab, set Provide a list of accessible networks to clients. Even with that set, certain cases such as Windows 10 may require additional changes to direct clients to send only specific traffic over the tunnel.

14.5.3 Scaling OpenVPN

Use IPsec Instead

As mentioned previously, where possible, use IPsec instead. IPsec is much more efficiently integrated into the operating system, and is capable of much greater throughput than OpenVPN.

Use UDP

UDP has less overhead for tunneled data, and if a client has to retransmit, it won’t compound the problem by retransmitting both inside and outside the tunnel. Unless there are extenuating circumstances which require TCP, use UDP.

Use TLS for Authentication Only

OpenVPN can use TLS for both authentication and for encryption of the control channel. Performing control channel encryption adds more overhead, which can add up with many clients. If control channel encryption is not required, consider using TLS for only authentication instead. No matter which option is chosen, traffic carried by OpenVPN is encrypted.

Encryption Algorithm

Use a CPU with AES-NI when possible, and use AES-GCM for the Encryption Algorithm when possible. Note that for AEAD ciphers such as AES-GCM, OpenVPN ignores the setting for Auth Digest Algorithm.

Note: AES-GCM can only be used in SSL/TLS mode, not Shared Key mode.

Use Negotiable Crypto Parameters

NCP can be used to set preferences so that more efficient ciphers can be preferred by clients where possible, but others can be used when necessary. Set high-priority selections such as *AES-128-GCM* first, followed by others like *AES-128-CBC*.

Split Tunneling

As mentioned in the general section above, split tunneling only sends traffic for specific subnets across the VPN rather than sending all traffic. With OpenVPN, this can be done by Unchecking the Redirect IPv4/IPv6 Gateway option(s) and configuring IPv4/IPv6 Local Network(s) entries instead. Clients may still override this behavior remotely, however, so check the client configurations as well.

Concurrent Connections

AZTCO-FW does not impose any connection limits by default, but an administrator may have chosen to configure a limit on the number of connections via the Concurrent Connections setting on servers. Ensure this is either unset or set high enough to accommodate the required number of users.

Disable Compression

Though using compression is tempting to squeeze extra throughput out of slower links, it is both inefficient and insecure. Most data sent across VPNs in modern environments is already encrypted or otherwise uncompressible, which wastes CPU when attempting to compress. Additionally, vulnerabilities such as VORACLE can allow attackers to glean information about encrypted data when it has been compressed. Disabling encryption will mitigate that attack and also reduce CPU overhead. On the server, set Compression to *Disable Compression*.

Duplicate Connections

Normally, if an OpenVPN client connects using the same username or certificate CN, the older connection is broken in favor of the new connection. This is more secure, but does not allow any given user to connect multiple times. Circumstances may necessitate supporting this, and in some environments it's not possible to give every device a unique username and/or certificate. Check Duplicate Connection in the OpenVPN server settings to allow multiple connections from the same user.

Topology

On recent versions of AZTCO-FW, OpenVPN defaults to *subnet* topology which uses addresses more efficiently, but if the VPN was configured initially on older versions, or if an older guide was followed, it may still be using *net30* topology. Using a common example tunnel network of 10.0.8.0/24, with *subnet* topology, the VPN can have a maximum of 252 users but with *net30*, it can only have 63. This is because in *net30* mode, each user receives a /30 subnet which utilizes four IP addresses for each user. In *subnet* mode, the server uses a single address and the client uses a single address, which is much more efficient.

Use UDP Fast I/O

This option is experimental but for those who have used it, it can result in much higher throughput. Not all platforms support it, however.

Increase Send/Receive Buffer

The default buffer size is safe, but not optimal. Increasing the buffer size to 512KiB on both sides can result in greater throughput. Results will vary by platform, internet link speed, and other factors. May require experimenting with multiple values to find the most efficient setting for a given environment.

Use Multiple Servers

OpenVPN is not multi-threaded so any single instance of OpenVPN is limited to using a single CPU. If a router has fast cores and not too many users, that may be OK, but it does not scale well. A workaround for this is to split users onto multiple servers. There are various means to reach this goal, including (but not limited to):

- Multiple servers on different WANs or ports, each with unique tunnel networks but otherwise identical settings (Same CA structure, encryption, etc).
  - Administrators could choose to manually configure pools of clients to connect to specific servers, but that does not scale well.
  - Clients may connect to any server configured in this manner so long as their settings line up properly.
  - Multiple servers can be listed in a single client configuration with additional remote statements.
  - Add remote-random to the client configuration so that clients will pick a random server when starting, which avoids overloading whichever server is listed first.
  - Servers could be run on multiple WANs to overcome single-circuit throughput limits.
- Multiple servers with completely unique settings (Different CA structure, different clients, etc)
  - More secure but more difficult to manage.
  - Clients must use different configurations to reach each server, no automated/built-in way to pick between them unless a specific client supports that function.
  - Good for isolating separate security levels (e.g. remote workers, remote administrators, vendors).

Process Efficiency

As a counterpoint to the above, each server will incur additional memory and other overhead to manage the process. When dealing with site-to-site VPNs, it is more efficient from a *memory* standpoint to use a single server with multiple clients (Peer to Peer SSL/TLS) vs servers for every node (Peer to Peer Shared Key, or SSL/TLS with a /30 tunnel network). If memory is a limiting factor, use fewer servers. If CPU overhead is the limiting factor, use separate servers.

14.6 OpenVPN

14.6.1 OpenVPN and IPv6

OpenVPN can connect a site-to-site tunnel to either an IPv4 address or an IPv6 address and both IPv4 and IPv6 traffic may be passed inside of an OpenVPN tunnel at the same time. IPv6 is supported both in site-to-site and mobile clients, and it can be used to deliver IPv6 to a site that only has IPv4 connectivity. In order to ensure mobile client support for IPv6, obtain the client software from the OpenVPN client export package, or download a client based on OpenVPN
2.3 or newer.

14.6.2 OpenVPN Configuration Options

This section describes all of the available options with OpenVPN and when they are typically used. Subsequent sections cover examples of configuring site-to- site and remote access VPNs with OpenVPN, using the most common options and a minimal configuration.

Server Configuration Options

These options are available in one or more modes for OpenVPN server instances, managed from VPN > OpenVPN, on the Servers tab.

Disable this server

Check this box and click Save to retain the configuration, but not enable the server. The process for this instance will be stopped, and all peers/clients will be disconnected from this server. Any other active servers are unaffected.

Server Mode

This is the role for the server, which specifies how routers or users will connect to this server instance. Changing this will also affect what options will appear on the rest of the page, so only relevant choices are displayed.

Peer to Peer (SSL/TLS) A connection between local and remote networks that is secured by SSL/TLS. This choice offers increased security as well as the ability for the server to push configuration commands to the remote peer router when using a 1:many style setup. Remote peer routers can also have certificates revoked to remove access if they become compromised.

Peer to Peer (Shared Key) A connection between local and remote networks that is secured by a single Shared Key configured on both nodes. This choice is easier to setup, but is less secure. If a shared key is compromised, a new key must be generated and then copied to any router or client using the old shared key. In this mode, a separate server instance is needed for each client.

Remote Access (SSL/TLS) This choice is a mobile client setup with per-user X.509 certificates. As with the peer-to-peer SSL/TLS connection type, using this method offers increased security as well as the ability for the server to push configuration commands to clients. Mobile clients can also have keys revoked to remove access if a key is compromised, such as a stolen or misplaced laptop.

Remote Access (User Auth) A client access server that does not use certificates, but does require the end user to supply a username and password when making a connection. This is not recommended unless authentication is handled externally by LDAP or RADIUS.

Remote Access (SSL/TLS + User Auth) The most secure choice offered. Not only does it get the benefits of other SSL/TLS choices, but it also requires a username and password from the client when it connects. Client access can be removed not only by revoking the certificate, but also by changing the password. Also, if a compromised key is not immediately discovered, the danger is lessened because it is unlikely that the attacker has the keys and the password. When using the OpenVPN wizard, this is the mode which is configured during that process.

Protocol

TCP or UDP may be selected, or their IPv6-enabled counterparts, TCP6 or UDP6. An OpenVPN server instance can currently only bind to either IPv4 or IPv6, but not both at the same time. UDP is the most reliable and fastest choice for running OpenVPN, and it should always be used when possible. In some rare cases TCP can be used to work around limitations, such as bypassing some firewalls by running an OpenVPN server on TCP port 443.

Connectionless protocols such as UDP are always preferable when tunneling traffic. TCP is connection oriented with guaranteed delivery, so any lost packets are retransmitted. This sounds like a good idea on the surface but TCP retransmissions will cause performance to degrade significantly on heavily loaded Internet connections or those with consistent packet loss.

TCP traffic frequently exists within tunnels and it is undesirable to retransmit lost packets of encapsulated VPN traffic. In cases where TCP is wrapped around TCP, such as a VPN tunnel using TCP as a transport protocol, when a packet is lost both the outer and inner lost TCP packets will be re-transmitted. Infrequent occurrences of this will be unnoticeable but recurring loss will cause significantly lower performance than UDP. If the traffic inside the tunnel requires reliable delivery, it will be using a protocol such as TCP which ensures that and will handle its own retransmissions.

Device Mode

OpenVPN can run in one of two device modes: *tun* or *tap*:

tun Works on OSI layer 3 and performs routing on point-to-point interfaces. tap Can

work at OSI layer 2 and can perform both routing and bridging if necessary.

Note: Not all clients support tap mode, using tun is more stable and more widely supported. Specifically, clients such as those found on Android and iOS only support tun mode in the Apps most people can use. Some Android and iOS OpenVPN apps that require rooting or jailbreaking a device do support tap, but the consequences of doing so can be a bit too high for most users.

Interface

Selects the interface, VIP, or failover group that the OpenVPN server instance will listen upon for incoming connections. This also controls which interface the traffic from the server will exit.

Several types of options are listed in the drop-down for Interface, and some have special behavior or use cases:

Interfaces OpenVPN will bind to the interface address. If the interface is dynamic, such as DHCP, OpenVPN will automatically bind to the new address if it changes.

VIPs OpenVPN will bind only to the specified VIP (IP Alias or CARP type)

Gateway Groups For use with failover groups, OpenVPN will bind to the address of the interface that is currently active in the group. If that interface gateway becomes unreachable, the next one will be used instead, and so on.

Localhost Useful for Multi-WAN deployments, binding to localhost and utilizing port forwards to accept connections from several interfaces and/or ports is a versatile way to provide redundant OpenVPN connectivity for connecting clients.

Any Binds to every address on every interface. Though tempting, this option is not recommended. When used with UDP, replies to Internet clients will always exit back out the default gateway WAN, which may be undesirable.

Local port

The local port is the port number OpenVPN will use to listen. Firewall rules need to allow traffic to this port and it must be specified in the client configuration. The port for each server must be unique for each interface.

Description

Enter a description for this server configuration, for reference.

Cryptographic Settings

This section controls how traffic to and from clients is encrypted and validated.

Shared Key

When using a shared key instance, either check the Automatically generate a shared key box to make a new key, or uncheck the box to paste in a shared key from an existing OpenVPN tunnel. When generating the key automatically, return to the edit screen for this tunnel later to obtain the key which may be copied to the remote router.

TLS Authentication

TLS, or Transport Layer Security, provides session authentication to ensure the validity of both the client and the server. Check the box to Enable authentication of TLS packets if desired. If there is no existing TLS key, leave Automatically generate a shared TLS authentication key checked. If key already exists, uncheck that option and then paste it into the provided entry box. When generating the key automatically, return to the edit screen for this tunnel later to obtain the key which may be copied to the remote router or client.

Warning: When using an SSL/TLS mode, we strongly recommend using TLS Authentication as well. In addition to the added security benefit from the key requirement, a TLS key also helps protect against some SSL-based attacks such as Heartbleed.

Peer Certificate Authority

Select the certificate authority used to sign the client or peer certificate(s) for this OpenVPN server instance. If none appear in this list, first import or generate a certificate authority under System > Cert Manager, on the CAs tab.

Peer Certificate Revocation List

This optional field is for the Certificate Revocation List (CRL) to be used by this tunnel. A CRL is a list of certificates made from a given CA that are no longer considered valid. This could be due to a certificate being compromised or lost, such as from a stolen laptop, spyware infection, etc. A CRL can be created or managed from System > Cert Manager, on the Certificate Revocation tab.

Server Certificate

A server certificate must be chosen for each OpenVPN server instance. If none appear in this list, first import or generate a certificate authority under System > Cert Manager, on the Certificates tab.

DH Parameters Length

The Diffie-Hellman (DH) key exchange parameters are used for establishing a secure communications channel. They may be regenerated at any time, and are not specific to an OpenVPN instance. That is, when importing an existing OpenVPN configuration these parameters do not need to be copied from the previous server. The length of the desired DH parameters may be chosen from the drop-down box, either 1024, 2048, or 4096.

Note: Due to the heavy computation involved in generating DH keys, a pre- generated set for each key type is used. New DH parameters may be generated manually by using the following shell commands:

```
# /usr/bin/openssl dhparam 1024 > /etc/dh-parameters.1024
# /usr/bin/openssl dhparam 2048 > /etc/dh-parameters.2048
# /usr/bin/openssl dhparam 4096 > /etc/dh-parameters.4096
```

Encryption algorithm

The cryptographic cipher to be used for this connection. The default is *AES- 128-CBC*, which is AES 128 bit Cipher Block Chaining. This is a fine choice for most scenarios.

See also:

Hardware Crypto for more information on using cryptographic accelerators and choosing an encryption algorithm.

Auth Digest Algorithm

Selects the message digest algorithm to use for HMAC authentication of incoming packets.

Note: OpenVPN defaults to SHA1 when this option is not specified, so unless both sides are set to a known value, use SHA1 here.

Hardware Crypto

If available, this option controls which hardware cryptographic accelerator will be used by OpenVPN. When left unspecified, OpenVPN will choose automatically based on what is available in the Operating System.

If this firewall device has a hardware cryptographic accelerator, choose BSD Cryptodev Engine, or select the specific device if it appears in the list. Most accelerator boards use the BSD cryptodev engine, so when in doubt, select that. This setting will allow OpenVPN to take advantage of the hardware acceleration. An encryption algorithm supported by the accelerator must also be selected. Refer to the hardware documentation for information on ciphers supported by the accelerator.

Certificate Depth

This option limits the length of a certificate chain before it fails validation. This defaults to *One (Client+Server)* so that if somehow an unauthorized intermediate CA is generated, certificates signed by the rogue intermediate would fail validation. In cases when chaining with intermediates is required, this limit can be raised.

Strict User-CN Matching

For SSL/TLS+User Authentication server, when enabled, this option enforces a match between the username supplied by the user and the Common Name of their user certificate. If the two do not match, the connection is rejected. This prevents users from using their own credentials with another person’s certificate and vice versa.

Tunnel Settings

The tunnel settings section governs how traffic flows between the server and clients, including routing and compression.

IPv4/IPv6 Tunnel Network

These are the pools of addresses to be assigned to clients upon connecting. The server’s end of the OpenVPN configuration will use the first address in this pool for its end of the connection, and assign additional addresses to connected clients as needed. These addresses are used for direct communication between tunnel endpoints, even when connecting two existing remote networks. Any subnet may be chosen provided that it is not in use locally or at any remote site. One or both of IPv4 Tunnel Network and IPv6 Tunnel Network may be entered, or in the case of a tap bridge, neither.

Warning: Currently, limitations in OpenVPN itself prevent running with only an IPv6 Tunnel Network configured. When an IPv6 Tunnel network is defined, an IPv4 Tunnel Network must also be specified, even if it is not used.

For a site-to-site SSL/TLS server using IPv4, the IPv4 Tunnel Network size can alter how the server behaves. If x.x.x.x/30 is entered for the IPv4 Tunnel Network then the server will use a peer-to-peer mode much like Shared Key operates: It can only have one client, does not require client-specific overrides or *iroutes*, but also cannot push routes or settings to clients. If an IPv4 Tunnel Network larger than that is used, such as x.x.x.x/24, the server will accept multiple clients and can push settings, but does require *iroutes*.

Bridging Options

When using *tap* mode, additional options are shown that control bridging behavior in OpenVPN and client address assignment. These are covered in *Bridging OpenVPN Connections to Local Networks*

Redirect Gateway

When the Redirect Gateway option is selected the server will push a message to clients instructing them to forward *all* traffic, including Internet traffic, over the VPN tunnel. This only works in SSL/TLS modes with a tunnel network larger than a /30 subnet.

IPv4/IPv6 Local network

These fields specify which local networks are reachable by VPN clients, if any. A route for these networks is pushed to clients connecting to this server. If multiple routes for subnets of a particular family are needed, enter the subnets separated by a comma, e.g. 192.168.2.0/24, 192.168.56.0/24.

This function relies upon the ability to push routes to the client, so for IPv4 it is only valid in an SSL/TLS context when a tunnel network larger than a /30 is in use. It will always work for IPv6 provided a similar too-small mask isn't set.

IPv4/IPv6 Remote Network

This option only appears when a Peer-to-Peer type connection is used, and is not available for mobile clients. Routes table entries are added to the firewall for the specified subnets, which hand the traffic over to this OpenVPN instance for processing. If more than one Remote network subnet is needed, enter the subnets separated by a comma, e.g. 192.168.2.0/24, 192.168.56.0/24.

Concurrent Connections

Specifies the number of clients that may be simultaneously connected to this OpenVPN server instance at any given time. This is a collective limit for all connected clients, not a per-user setting.

Compression

When compression is enabled, traffic crossing the OpenVPN connection will be compressed before being encrypted. This saves on bandwidth usage for many types of traffic at the expense of increased CPU utilization on both the server and client. Generally this impact is minimal, and enabling compression is beneficial for nearly any usage of OpenVPN over the Internet.

For high speed connections, such as the usage of OpenVPN across a LAN, high speed low/latency WAN, or local wireless network, this may be undesirable, as the delay added by the compression may be more than the delay saved in transmitting the traffic. If nearly all of the traffic crossing the OpenVPN connection is already encrypted (such as SSH, SCP, HTTPS, among many other protocols), do not enable LZO compression because encrypted data is not compressible and the LZO compression will cause slightly more data to be transferred than would be without compression. The same is true if the VPN traffic is almost entirely data that is already compressed.

This selector controls the handling of LZO compression for this OpenVPN instance. There are four possible settings each with slightly different behavior.

No Preference Omits the compression directives from the OpenVPN configuration entirely. No compression will be performed, but this may be overridden by other methods such as Client-Specific overrides or advanced options.

Disabled - No Compression Explicitly disables compression in the configuration

Enabled with Adaptive Compression Enables compression with a periodic test to ensure the traffic is able to be compressed. If compression is not optimal, it will be disabled until it is tested again. This option strikes the best balance since it will compress data when it will help, but does not compress data when it is hindering performance.

Enabled without Adaptive Compression Explicitly enables compression to be on at all times without testing the traffic.

Type-of-Service

When this option is enabled OpenVPN will set the Type-of-Service (TOS) IP header value of tunnel packets to match the encapsulated packet value. This may cause some important traffic to be handled faster over the tunnel by intermediate hops, at the cost of some minor information disclosure.

The most common example is VoIP or video traffic. If the TOS bit is set to reflect the priority of the traffic it can help QoS along the path, but someone intercepting the traffic could see the TOS bit and gain some knowledge about the contents of the traffic inside the tunnel. For those who rely on TOS bits for QoS, the benefit may outweigh the information leak.

Inter-Client Communication

This option controls whether or not connected clients are able to communicate with one another. To allow this behavior, check the option. When unchecked, clients can only send traffic to the server or destinations beyond the server such as routed networks or the Internet.

Typically in remote access style deployments it is unnecessary for clients to reach each other, but there are some corner cases when it can be helpful. One example is remote web developers working together and running test servers on their local systems. With this option activated, they can reach the other test serves for collaborative development.

Duplicate Connections

By default OpenVPN will associate an IP address from its tunnel network with a specific certificate or username for a given session. If the same certificate connects again, it would be assigned the same IP address and either disconnect the first client or cause an IP conflict where neither client will receive proper data. This is primarily for security reasons so the same certificate cannot be used by multiple people simultaneously. We recommend a unique certificate be used for each connecting user. Otherwise if a client is compromised there is no way to revoke that one client alone, certificates would need to be reissued to all clients that share the same certificate.

If a setup that uses the same certificate in multiple locations is an absolute requirement and cannot be avoided, check Duplicate Connections to allow the non-standard behavior of multiple clients using the same certificate or username.

Disable IPv6

When checked, IPv6 traffic forwarding is disabled for this OpenVPN instance.

Client Settings

These settings pertain to how clients connecting to this sever instance will behave.

Dynamic IP

Checking this box adds the *float* configuration option to the OpenVPN configuration. This allows clients to retain their connection if their IP address changes. similar to MOBIKE for IKEv2 in IPsec. For clients on Internet connections where the IP changes frequently, or mobile users who commonly move between different Internet connections, check this option to allow for stable connectivity. Where the client IP is static or rarely changes, not using this option offers a small security improvement.

Address Pool

When this option is enabled the server will assign virtual adapter IP addresses to clients from the subnet specified by the Tunnel Network option. When unchecked IP addresses will not be assigned automatically and clients will have to set their own static IP addresses manually in their client configuration files. Except in rare cases, this is almost always enabled.

Topology

By default OpenVPN on AZTCO-FW® software version 2.3 and later prefers a topology style of *subnet* when using a Device Mode of *tun*. This style allocates only one IP address per client rather than an isolated subnet per client. This is the only available style when using the *tap* Device Mode.

When the older *net30* topology for *tun* is chosen, OpenVPN allocates a /30 CIDR network (four IP addresses, two usable) to each connecting client. This style has a longer history, but can be confusing for administrators and users alike.

The Topology option is relevant only when supplying a virtual adapter IP address to clients using *tun* mode on IPv4. Some clients may require this even for IPv6, such as OpenVPN Connect, though in reality IPv6 always runs with a subnet topology even when IPv4 uses *net30*. OpenVPN version 2.1.3 or newer is required to use a *subnet* topology, and there were significant fixes to it in OpenVPN 2.3 as well, so using a current OpenVPN client version is important.

> Warning: The default in AZTCO-FW has been changed to *subnet* because the OpenVPN project has declared the *net30* style as deprecated, indicating it will be removed in future versions.
>
> Be aware, however, that some very old clients may break if this option is used, such as older versions of OpenVPN (Before 2.0.9, released nearly 10 years ago), Windows versions with older tun/tap drivers, or clients such as Yealink phones. Always make sure the client and associated drivers are fully up-to-date when using a *subnet* topology.

DNS Default Domain

When checked, a field will appear to specify the DNS domain name to be assigned to clients. To ensure name resolution works properly for hosts on the local network where DNS name resolution is used, specify the internal DNS domain name here. For Microsoft Active Directory environments, this would usually be the Active Directory domain name.

DNS servers

When checked, up to four DNS servers may be entered for use by the client while connected to the VPN. For Microsoft Active Directory environments, this is typically the Active Directory Domain Controllers or DNS servers for proper name resolution and authentication when connected via OpenVPN.

Force DNS Cache Update

When checked, this option will push a set of commands to Windows clients that will flush their DNS and restart caching to improve client handling of updated DNS servers from the VPN.

NTP servers

When checked, one or two NTP servers may be set for syncing clocks on clients. It can be an IP address or FQDN.

NetBIOS Options

When Enable NetBIOS over TCP/IP is checked, several other NetBIOS and WINS related options will appear. If the box is unchecked, these settings will be disabled.

Node Type

The NetBIOS node type controls how Windows systems will function when resolving NetBIOS names. It's usually fine to leave this to *none* to accept Windows' default.

The available options include:

b-node Use broadcasts for NetBIOS name resolution. This would not be used except in the case of a tap bridge.

p-node Point-to-point name queries to a WINS server. WINS has been mostly deprecated, so this option is not useful in modern Windows networks.

m-node Broadcast then query name server. Similar to b-node but will fall back to DNS.

h-node Query name server first, then use broadcast. This option is the most likely to succeed in a current network with proper, functional, DNS.

Scope ID

A NetBIOS Scope ID provides an extended naming service for NetBIOS over TCP/IP. The NetBIOS scope ID isolates NetBIOS traffic on a single network to only those nodes with the same NetBIOS scope ID.

WINS Servers

Checking this box allows two WINS servers to be defined which provides name resolution for clients accessing and browsing NetBIOS resources across the VPN. WINS has been largely deprecated and removed from use, so it's unlikely this will be needed in most modern environments.

Enable Custom Port

When checked, a non-default Management Port may be specified for use with the OpenVPNManage feature of the OpenVPN Client Export package. If multiple connections profiles are used on a single client using that interface, each must use a unique management port.

Custom options

While the AZTCO-FW web interface supports the most commonly used options, OpenVPN is very powerful and flexible and occasionally options that are unavailable in the web interface may be necessary. Such custom options may be added in using this entry box. These options are described further in *Custom configuration options*.

Verbosity level

Configures the amount of detail shown in the OpenVPN logs for this instance, useful for troubleshooting problems. Higher numbers will result in higher amounts of detail in the log. During normal operation the *default* selection is best.

Note: When set to higher levels, the OpenVPN status page and dashboard widget will cause additional logging as they interact with the Management process to poll information from the OpenVPN daemons.

Client Configuration Options

These options are available in one or more modes for OpenVPN client instances, managed from VPN > OpenVPN, on the Clients tab.

Many of these options are identical to the server options mentioned above, so only differences will be noted.

Server mode

For client instances, the server mode choices are limited to *Peer to Peer (SSL/TLS)* and *Peer to Peer (Shared Key)*, which pair with the server options of the same name and type.

Interface

This option selects the interface, VIP, or failover group that the OpenVPN client instance will use for outgoing connections.

When a CARP type VIP is selected for the Interface on OpenVPN Client instances, the OpenVPN instance will be stopped when the CARP VIP is in a backup state. This is done to prevent the secondary HA node from maintaining invalid routes or attempting to make outbound connections which can interfere with the active connection on the primary HA node.

Local Port

For clients, the local port is left blank in nearly every case so that a randomized local port will be used. This is more secure, but some implementations may require a specific source port. If a specific source port is required, fill it in as needed as needed.

Server host or address

The IP address or fully qualified domain name for the server.

Server Port

The port on which the server is listening, typically 1194

Proxy Settings

Proxy Host or Address The IP address or fully qualified domain name for a proxy server through which this client must connect.

Proxy Auth Extra Options Extra authentication options. When set to *basic* or *ntlm*, Username and Password fields are presented so that proxy authentication may be configured.

Server Hostname Resolution

When Infinitely Resolve Server is checked, the server host name will be resolved on each connection attempt. When unchecked, OpenVPN will only attempt to resolve it once. When using a hostname for the remote server address, this option should be checked.

User Authentication Settings

When using *Peer to Peer SSL/TLS* mode, a Username and Password may be specified in addition to, or instead of, a user certificate, depending on the requirements configured on the server.

Cryptographic Settings

The settings in this section are identical to those on their corresponding options on the server side except for the new Client Certificate option, where the certificate is selected for use by this client. This certificate (and the associated key, and CA Certificate) must be imported to this firewall before they can be chosen.

Shared Key / TLS Authentication

These options work similar to the server side counterparts, but be aware that the key from the server must be copied here, rather than generating a new key on the client.

Limit Outgoing Bandwidth

The value in this box, specified in bytes per second, is used to limit the speed of outgoing VPN traffic. When left blank, there is no limit. The value must be between *100* and *100000000*.

Don’t Pull Routes

When checked, the client will ignore routes pushed from the server. This is useful in cases when the server pushes a default gateway redirect when this client does not need one.

Don’t Add/Remove Routes

When checked, OpenVPN will not manage route table entries for this VPN. In this case, they must be managed manually. The routes that would normally be added are instead passed to --route-upscript using environmental variables.

14.6.3 Custom configuration options

OpenVPN offers dozens of configuration options, many beyond the most commonly used fields presented in the GUI. This is why the Advanced configuration box exists. Additional configuration options may be configured using this input area, separated by semicolons.

This section covers the most frequently used custom options individually. There are many more, though rarely needed. The OpenVPN man page details them all.

> Warning: Exercise caution when adding custom options, there is no input validation applied to ensure the validity of options used. If an option is used incorrectly, the OpenVPN client or server may not start. View the OpenVPN logs under Status > System logs on the OpenVPN tab to ensure the options used are valid. Any invalid options will result in a log message, followed by the option that caused the error:
>
> ```
> Options error: Unrecognized option or missing parameter(s)
> ```

Routing options

To add additional routes for a particular OpenVPN client or server, use the Local Network and Remote Network boxes as needed, using a comma- separated list of networks.

The route custom configuration option may also be used, but is no longer necessary. Some users prefer this method, however. The following example adds a route for 10.50.0.0/24:

```
route 10.50.0.0 255.255.255.0;
```

To add multiple routes, separate them with a semicolon:

```
route 10.50.0.0 255.255.255.0; route 10.254.0.0
255.255.255.0;
```

The route configuration option is used to add routes locally for networks that are reachable through the VPN. For an OpenVPN server configuration using PKI, additional routes may also be pushed to clients. The GUI can configure these using the Local Network field. To push the routes manually for 10.50.0.0/24 and 10.254.0.0/24 to all clients, use the following custom configuration option:

```
push "route 10.50.0.0 255.255.255.0"; push "route
10.254.0.0 255.255.255.0";
```

Redirecting the default gateway

OpenVPN also allows the default gateway to be redirected across the VPN, so all non-local traffic from the client is sent through the VPN. This is great for untrusted local networks such as wireless hotspots, as it provides protection against numerous attacks that are a risk on untrusted networks. This is configurable in the GUI now, using the Redirect Gateway checkbox in the OpenVPN instance configuration. To do this manually, add the following custom option:

```
push "redirect-gateway def1"
```

The same value may be used as a custom option on the client side by entering redirect-gateway def1 without specifying push . (Note the option is the letters “def” followed by the digit *one*, not the letter “L”.)

14.6.4 OpenVPN Firewall Rules

Permitting traffic to the OpenVPN server

A firewall rule must permit traffic to the OpenVPN server or clients will not be able to connect. Add a rule as follows:

- Navigate to Firewall > Rules, WAN tab
- Click ↧ to create a new rule at the top of the list
- Set Protocol to *UDP*
- Leave the Source set to *any*
- Set the Destination to *WAN Address*
- Set the Destination port to 1194 in this instance (or whichever port the server is using to listen)
- Enter a Description, such as Allow traffic to OpenVPN Server
- Click Save
- Click Apply changes

This rule is depicted in Figure *OpenVPN Server WAN Rule*.

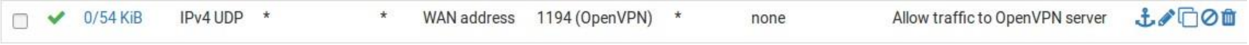

□	✓	0/54 KiB	IPv4 UDP	*	*	WAN address	1194 (OpenVPN)	*	none	Allow traffic to OpenVPN server	⚓✎❐⊘🗑

Fig. 1: OpenVPN Server WAN Rule

If the client source addresses are known and do not change, then the source of the rule could be altered to limit traffic from only those clients. This is more secure than leaving the server exposed to the entire Internet, but that is necessary to accommodate clients with dynamic IP addresses, roaming clients, and so on. The risk of leaving the service exposed with most OpenVPN configurations is minimal, especially in cases where TLS Authentication is employed. With certificate based authentication there is less risk of compromise than password- based solutions that are susceptible to brute forcing. This presumes a lack of security holes in OpenVPN itself, which to date has a solid security track record.

Allowing traffic over OpenVPN Tunnels

By default, all traffic is blocked from entering OpenVPN tunnels. To allow traffic from remote OpenVPN nodes to make connections to resources on the local side, firewall rules under Firewall > Rules, on the OpenVPN tab are required.

As with other aspects of the firewall, these rules will only match traffic coming into the firewall from the remote side, not traffic leaving from this side, so craft the rules accordingly. In cases when AZTCO-FW® software is on both ends and traffic is required to reach between local networks on both sides, then rules are required on both firewalls.

Add an OpenVPN rule which passes all traffic as follows:

- Navigate to Firewall > Rules, OpenVPN tab
- Click ↴ to create a new rule at the top of the list
- Set Protocol to *any*
- Enter a Description such as Allow all on OpenVPN
- Click Save
- Click Apply changes

To limit the traffic to only specific sources and destinations, adjust the rule(s) as needed. A strict ruleset is more secure, but more difficult to create.

Tip: Rules on the OpenVPN tab apply to all OpenVPN server and client instances. The OpenVPN interface may also be assigned (*Assigning OpenVPN Interfaces*) in which case there will be a separate firewall rule tab for that VPN, upon which rules can pass traffic for that specific VPN.

14.6.5 OpenVPN clients and Internet Access

For OpenVPN Remote Access clients to reach the Internet through the OpenVPN connection, Outbound NAT is required to translate their traffic to the WAN IP address of the firewall. The default Automatic Outbound NAT rules cover this, but if Manual Outbound NAT is in use, manual rules are necessary to perform outbound NAT on traffic from sources that include the OpenVPN tunnel network or remote network(s).

See also:

Outbound NAT for more details on Outbound NAT.

14.6.6 Assigning OpenVPN Interfaces

In order to do complex NAT, policy routing, or tunnel-specific filtering, the OpenVPN interface must be assigned as an OPT interface and configured accordingly.

Assigning the OpenVPN interface enables several beneficial changes for advanced control of VPN traffic:

- Adds a firewall tab under Firewall > Rules
- Adds *reply-to* to rules on the VPN interface tab to help with return routing
- Adds a Gateway entry for the far side of the VPN for policy routing
- Allows the interface to be selected elsewhere in the GUI and packages
- Allows more fine-grained control of Port Forwards and Outbound NAT for the VPN

Interface assignment and configuration

- Navigate to Interfaces > Assignments
- Select the appropriate ovpns or ovpnc interface in Available network ports, the description of the VPN is printed for reference.
- Click + Add to assign the interface as a new OPT interface (e.g. OPT1) Figure *Assign OpenVPN Interface* shows ovpns1 assigned as OPT1.

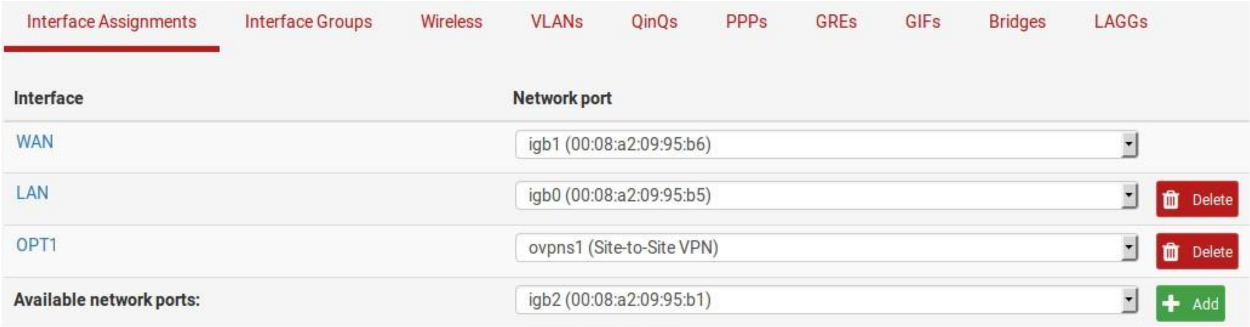

Interface	Network port	
WAN	igb1 (00:08:a2:09:95:b6)	
LAN	igb0 (00:08:a2:09:95:b5)	Delete
OPT1	ovpns1 (Site-to-Site VPN)	Delete
Available network ports:	igb2 (00:08:a2:09:95:b1)	+ Add

Fig. 2: Assign OpenVPN Interface

- Navigate to the Interface configuration page, Interfaces > OPTx
- Check Enable
- Enter an appropriate Description which will become the interface name (e.g. VPNServer)
- Select *none* for both IPv4 Configuration Type and for IPv6 Configuration Type

> Note: This will not configure any IP address information on the interface, which is necessary since OpenVPN itself must configure these settings.

- Click Save

- Click Apply Changes

This does not change the functionality of OpenVPN, it makes the interface available for firewall rule, NAT, and gateway purposes, among other uses.

After assigning the OpenVPN interface, edit the OpenVPN server or client and click Save once there as well to reinitialize the VPN. This is necessary for the VPN to recover from the assignment process.

Filtering with OpenVPN

When the OpenVPN interface is assigned, a tab is present under Firewall > Rules dedicated to only this single VPN. These rules govern traffic coming in from the remote side of the VPN and they even get the pf reply-to keyword which ensures traffic entering this VPN interface will exit back out the same interface. This can help with some more advanced NAT and configuration scenarios.

Note: Rules added here are processed *after* the OpenVPN tab rules, which are checked first. In order to match the rules on an assigned VPN tab, the traffic must not match any rules on the OpenVPN tab. Remove any "Allow All" style rules from the OpenVPN tab and craft more specific rules instead.

See also:

For more information on firewall rules, refer to *Firewall*.

Policy Routing with OpenVPN

When the OpenVPN interface is assigned and enabled, an automatic gateway entry is added under System > Routing, on the Gateways tab. With this, traffic can be directed into the VPN using the Gateway field on LAN or other internal interface firewall rules.

When used with a VPN to reach Internet sites, more configuration may be required. Either outbound NAT must be performed on the VPN interface before it leaves (for VPN services such as PIA, StrongVPN and similar) or the NAT must be done on the other side before it reaches the actual Internet connection.

See also:

See *Policy routing* for more information on policy routing.

Warning: Do not use this automatic gateway for static routes. Use the Remote Network field in the VPN configuration. Defining a static route using the automatic OpenVPN gateway will not work properly.

NAT with OpenVPN

When the OpenVPN interface is assigned NAT rules can also be applied the same as with any other interface. This is useful when connecting two conflicting subnets or for making NAT rules specific to this one VPN connection (outbound NAT, port forwards, or 1:1 NAT)

14.6.7 OpenVPN and Multi-WAN

OpenVPN is multi-WAN capable, with some caveats in certain circumstances. This section covers multi-WAN considerations with OpenVPN server and client configurations.

OpenVPN assigned to a Gateway Group

A Gateway Group (*Gateway Groups*) may be selected as the Interface for an OpenVPN instance. Such a gateway group must be configured for failover only, not load balancing. Failover groups only have one gateway per tier. When creating the gateway group, a VIP may also be chosen for use with a specific gateway. When selected for a VPN server, the interface or VIP of the Tier 1 gateway in the group will be used first. If that gateway goes down, it will move to tier 2, and so on. If the tier 1 gateway comes back up, the VPN will resume operating on that WAN immediately. When used for a VPN server, this means that the server is only active on one WAN at a time. Some of the other methods described below may be better for most common circumstances, such as needing both WANs usable concurrently with the VPN. When used with OpenVPN clients, the outbound interface will be switched according to the gateway group tiers.

OpenVPN servers and multi-WAN

OpenVPN servers can be used with any WAN connection, though the means of doing so will vary depending on the specifics of a given configuration.

OpenVPN server using TCP

TCP is not the preferred protocol for OpenVPN. However, using TCP can make multi-WAN OpenVPN easier to configure when the VPN is using an interface setting of *any*. OpenVPN servers using TCP will work properly on all WANs where the firewall rules allow the traffic to the OpenVPN server. A firewall rule is required for each WAN interface. This method should be considered a last resort, and only used if the other methods are not viable.

Note: This works because of the connection-oriented nature of TCP. The OpenVPN can reply back to the other end with the proper source preserved since it is part of an open connection.

OpenVPN server using UDP

OpenVPN servers with UDP are also multi-WAN capable, but with some caveats that aren’t applicable with TCP.

These OpenVPN limitations are due to the connectionless nature of UDP. The OpenVPN instance replies back to the client, but the Operating System selects the route and source address based on what the routing table believes is the best path to reach the other side. For non-default WANs, that will not be the correct path.

Multiple Server Method

In some cases, each WAN must have its own OpenVPN server. The same certificates may be for all the servers. Only two parts of the OpenVPN configuration must change:

Tunnel Network Each server must have a unique Tunnel Network that does not overlap with any other tunnel network or internal subnet.

Interface Each OpenVPN server must specify a different WAN Interface.

Port forward method

An easier and more flexible option is to bind the OpenVPN server to the *LAN* interface or *Localhost* and use a port forward from each WAN to direct the OpenVPN port to the service. Using this method the reply-to functionality in pf will ensure that the return traffic flows back to the proper source via the intended interface.

This method requires some minor manual intervention when used with the client export package. The Host Name Resolution option must be set to one of the automatic port forward methods otherwise the default export settings would leave it attempting to connect to the wrong address. See *OpenVPN Client Export Package* for details

Automatic Failover for Clients

Multiple remote servers can be configured on OpenVPN clients. If the first server cannot be reached, the second will be used. This can be used in combination with a multi-WAN OpenVPN server deployment to provide automatic failover for clients. If the OpenVPN servers are running on IP addresses 198.51.100.3 and 203.0.113.5, both using port 1194, the remote lines in the client configuration file will be as follows:

```
remote 198.51.100.3 1194 udp remote
203.0.113.5 1194 udp
```

For clients configured on AZTCO-FW® software, the first remote is configured by the Server Host or Address* field in the GUI. The second ``remote`` is specified in the **Advanced field.

This method has three notable behaviors that some may find undesirable:

- It will take at least 60 seconds to detect a failure and switch to the next server.
- Any connection failure will cause it to try the second server, even if it is not a WAN failure.
- It will not “fail-back”. Once a client connects to the second server IP address it will stay there until disconnected.

OpenVPN Clients and Multi-WAN

To use an OPT WAN interface, select it as the Interface. OpenVPN clients configured on the firewall will respect the chosen Interface and a static route is added automatically behind the scenes to ensure traffic takes the correct path.

If the interface is instead set to *any*, the client will follow the system routing table when making the connection to the OpenVPN server. In this case a manual static route will be required to direct traffic to the remote endpoint over the desired WAN.

OpenVPN Site-to-Site with Multi-WAN and OSPF

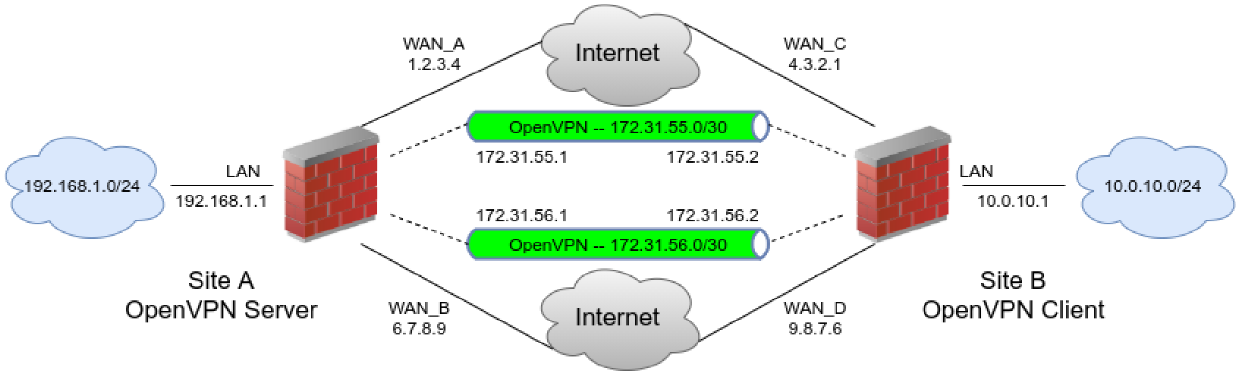

Fig. 3: Example OpenVPN Setup Involving OSPF Across Multiple WANs

Building upon concepts from earlier in the chapter, it is possible to configure a redundant VPN using a dynamic routing protocol such as OSPF as seen in Figure *Example OpenVPN Setup Involving OSPF Across Multiple WANs*.

First, setup shared key site-to-site OpenVPN instances on each WAN for the remote sites. Do not fill in the Remote Networks fields on either side, only Tunnel Network addresses.

- Setup two servers on the local side, each on a different port. Use two distinct, non-overlapping tunnel networks

  (e.g. 172.31.55.0/30 and 172.31.56.0/30)
- Setup two clients on the remote firewall, each paired up with one of the above servers, matching the IP addresses and port numbers involved.
- Ensure the clients are set for their specific WAN, choose the interface from the drop-down menu, or a CARP VIP that is on one of the WANs being used.
- Ensure these OpenVPN connections link up between client and server. The tunnel address on both sides will respond to a ping when they are working correctly. If the tunnels do not establish, see *Troubleshooting OpenVPN* for suggestions on troubleshooting the connection.
- Ensure the OpenVPN firewall rules allow all traffic or at least allow OSPF traffic from a source of the tunnel networks to a destination of any. The destination on the traffic will be a multicast address, which can be used to filter specifically if needed, but there isn't much to be gained in the way of security if the source is locked down in the rules as the traffic cannot leave that segment.

Once both instances are connected, configure OSPF.

- Install the Quagga_OSPF package from System > Packages, Available Packages tab on both firewalls.
- Navigate to Services > Quagga OSPFd, Interfaces tab
- Add each OpenVPN interface
  - Set the cost to 10 on the primary link and 20 on the secondary, and so on
  - Add the LAN and other internal interfaces as passive interfaces
- Navigate to the Global Settings tab
- Enter a Master Password. It doesn't matter what it's set to, it is used internally for accessing the status daemon.

- Set the Router ID to an IP-address-like value, (e.g. 10.3.0.1.) The Router ID is unique on each device, which is why setting it to the LAN IP address of a router is a good practice.
- Set the Area ID which is also an IP-address-like value. The Area ID is typically set to 0.0.0.0 or 0.0.0.1, but any properly-formatted value may be used. The Area ID is the same for all routers involved in this VPN
- Click Save

Once OSPF has been configured on all routers, they will attempt to form a neighbor relationship.

After OSPF has been setup on both ends the Status tab will show a full peering with each instance on each wan if they connected properly, and the routes obtained via OSPF will be listed. Once that happens, try unplugging/replugging WANs and refreshing the status while running some test traffic across the VPN, such as an ICMP ping.

14.6.8 OpenVPN and CARP

OpenVPN works well with High Availability using CARP. To provide a high availability OpenVPN solution with CARP, configure the OpenVPN server or client to use the CARP VIP with the Interface option and configure clients to connect to that CARP VIP.

When XMLRPC Configuration Synchronization settings are enabled, OpenVPN instances will automatically synchronize. The connection state isn't retained between hosts so clients must reconnect after failover occurs, but OpenVPN will detect the connection failure and reconnect within a minute or so of failover. High Availability and CARP are discussed further in *High Availability*.

When a CARP VIP is selected as the Interface for an OpenVPN instance the firewall will automatically shut down OpenVPN client instances as needed when a CARP node is in a BACKUP state. This prevents OpenVPN from making unnecessary outbound connections in client mode. When the CARP VIP status transitions to MASTER, the OpenVPN instances are started automatically.

14.6.9 Sharing a Port with OpenVPN and a Web Server

To be extra sneaky or careful with an OpenVPN server, take advantage of the port-share capability in OpenVPN that allows it to pass any non-OpenVPN traffic to another IP address behind the firewall. The usual use case for this would be to run the OpenVPN server on port tcp/443 while letting OpenVPN hand off the HTTPS traffic to a web server in place of a port forward.

Often on locked-down networks, only ports like 80 and 443 will be allowed out for security reasons and running OpenVPN instances on these allowed ports can help users get out in situations where access may otherwise be restricted.

To set this up, configure an OpenVPN server to listen on TCP port 443 and add a firewall rule to pass traffic to the WAN IP address or VIP used for OpenVPN on port 443. No additional port forwards or firewall rules are necessary to pass the traffic to the internal IP.

In the custom options of the OpenVPN instance, add the following:

```
port-share x.x.x.x 443
```

Where x.x.x.x is the internal IP address of the web server to which the non-VPN traffic will be forwarded by OpenVPN.

Now if an OpenVPN client is pointed to the public address it will connect to the VPN, and if a web browser is pointed at the same IP address, it will be connected to the web server.

14.6.10 Controlling Client Parameters via RADIUS

When using RADIUS as an authentication source for a VPN, AZTCO-FW® software supports receiving some client configuration parameters from the RADIUS server as reply attributes. The following values may be specified:

Cisco-AVPair inacl= Inbound firewall rules to govern traffic from the client to the server. Given in Ciscostyle ACL format (e.g. permit tcp any any) subnet masks are specified wildcard style.

Cisco-AVPair outacl= Outbound firewall rules to govern traffic from the server to the client. Formatted the same as the inacl parameter.

Cisco-AVPair dns-servers= DNS servers to push to the client. Multiple servers may be specified, separated by spaces.

Cisco-AVPair route= Additional route statements to push to the client. Specified as x.x.x.x y.y. y.y where the first parameter is a network address and the second is a subnet mask.

Framed-IP-Address= The IP address to assign to the client. When using a *subnet* style Topology the RADIUS server must also send back a Framed-Mask set appropriately for the Tunnel Network of the VPN. When using a *net30* style Topology, the client receives this IP address and the server side is set as one IP address lower than the address given to the client.

14.6.11 OpenVPN Adapter Address ICMP Behavior

Sometimes OpenVPN will not respond to ping on certain virtual addresses used solely for routing endpoints when using the *net30* topology. Do not rely on pinging the OpenVPN endpoint addresses as a means of determining if the tunnel is passing traffic properly. Instead, ping something in the remote subnet, such as the LAN IP of the server.

According to the OpenVPN FAQ, in the section titled *Why does OpenVPN's "ifconfig-pool" option use a /30 subnet (4 private IP addresses per client) when used in TUN mode?*:

> As 192.168.1.5 is only a virtual IP address inside the OpenVPN server, used as an endpoint for routes, OpenVPN doesn't bother to answer pings on this address, while the 192.168.1.1 is a real IP address in the servers O/S, so it will reply to pings.

This may seem a little counter-intuitive, since on the server something like this is seen in the ifconfig output:

```
tun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> metric 0 mtu 1500 inet6
          fe80::202:b3ff:fe03:8028%tun0 prefixlen 64 scopeid 0xc inet 192.168.100.1 --> 192.168.100.2
          netmask 0xffffffff
          Opened by PID 27841
```

While the client shows:

```
tun0: flags=8051<UP,POINTOPOINT,RUNNING,MULTICAST> metric 0 mtu 1500 inet6
          fe80::202:b3ff:fe24:978c%tun0 prefixlen 64 scopeid 0xa inet 192.168.100.6 --> 192.168.100.5
          netmask 0xffffffff
          Opened by PID 1949
```

In this case, *.5* or *.1* will not likely respond to ping. *.5* because it's a virtual address, and *.1* because the client has no route to reach it directly. The *.5* and *.6* addresses are part of a */30* that goes from *.4* to *.7*, and trying to ping *.1* would go out the default route instead.

There are many cases where the far side of an OpenVPN tunnel can respond to ping, but not the local. This is also counter-intuitive, but works especially in cases where there is a site-to-site link. If the server shows its tun addresses as "x.x.x.1 -> x.x.x.2" and the client shows the reverse - "x.x.x.2 -> x.x.x.1", then the far side will likely respond to ping from both ends.

In contrast, when using "topology subnet" these virtual addresses and /30 subnets are not used so these issues are not present.

See also:

- *OpenVPN Client Export Package*
- *Checking the Status of OpenVPN Clients and Servers*
- *OpenVPN Logs*
- *Connecting OpenVPN Sites with Conflicting IP Subnets*
- *OpenVPN Remote Access Configuration Example*
- *Authenticating OpenVPN Users with FreeRADIUS*
- *Authenticating OpenVPN Users with RADIUS via Active Directory*
- *Installing OpenVPN Remote Access Clients*
- *Installing the OpenVPN Client on iOS*
- *Adding OpenVPN Remote Access Users*
- *OpenVPN Site-to-Site Configuration Example with Shared Key*
- *Routing Internet Traffic Through A Site-To-Site OpenVPN Tunnel*
- *OpenVPN Site-to-Site Configuration Example with SSL/TLS*
- *Configuring a Single Multi-Purpose OpenVPN Instance*
- *Bridging OpenVPN Connections to Local Networks*
- *Connecting to an OpenVPN Access Server*
- *Troubleshooting OpenVPN*
- *Troubleshooting OpenVPN Internal Routing (iroute)*
- *Troubleshooting OpenVPN Push Routes*
- *Troubleshooting OpenVPN Remote Access Client IP Address Assignments*
- *Troubleshooting Windows OpenVPN Client Connectivity*
- *Troubleshooting Windows/SMB Share Access from OpenVPN Clients*

OpenVPN is an open source SSL VPN solution that can be used for remote access clients and site-to-site connectivity. OpenVPN supports clients on a wide range of operating systems including all the BSDs, Linux, Android, Mac OS X, iOS, Solaris, Windows 2000 and newer, and even some VoIP handsets.

Every OpenVPN connection, whether remote access or site-to-site, consists of a server and a client. In the case of site-to-site VPNs, one firewall acts as the server and the other as the client. It does not matter which firewall possesses these roles. Typically the location of the primary firewall will provide server connectivity for all remote locations, whose firewalls are configured as clients. This is functionally equivalent to the opposite configuration the primary location configured as a client connecting to servers running on the firewalls at the remote locations. In practice, the servers are nearly always run on a central location.

There are several types of authentication methods that can be used with OpenVPN: shared key, X.509 (also known as SSL/TLS or PKI), user authentication via local, LDAP, and RADIUS, or a combination of X.509 and user authentication. For shared key, a single key is generated that will be used on both sides. SSL/TLS involves using a trusted set of certificates and keys. User authentication can be configured with or without SSL/TLS, but its use is preferable where possible due to the increased security is offers.

The settings for an OpenVPN instance are covered in this chapter as well as a run-through of the OpenVPN Remote Access Server wizard, client configurations, and examples of multiple site-to-site connection scenarios.

For general discussion of the various types of VPNs available in AZTCO-FW and their pros and cons, see *Virtual Private Networks*.

14.6.12 OpenVPN and Certificates

Using certificates is the preferred means of running remote access VPNs, because it allows access to be revoked for individual machines. With shared keys, either a unique server and port for must be created for each client, or the same key must be distributed to all clients. The former gets to be a management nightmare, and the latter is problematic in the case of a compromised key. If a client machine is compromised, stolen, or lost, or otherwise needs revoked, the shared key must be re-issued to all clients. With a PKI deployment, if a client is compromised, or access needs to be revoked for any other reason, simply revoke that client's certificate. No other clients are affected.

The AZTCO-FW GUI includes a certificate management interface that is fully integrated with OpenVPN. Certificate authorities (CAs) and server certificates are managed in the Certificate Manager in the web interface, located at System > Cert Manager. User certificates are also managed in the web interface, as a part of the built-in user manager found at System > User Manager. Certificates may be generated for any user account created locally on the firewall except for the default admin account. For further information on creating a certificate authority, certificates, and certificate revocation lists, see *Certificate Management*.

14.7 IPsec

14.7.1 IPsec Configuration

IPsec offers numerous configuration options, affecting the performance and security of IPsec connections. Realistically, for low to moderate bandwidth usage it matters little which options are chosen here as long as DES is not used, and a strong pre-shared key is defined, unless the traffic being protected is so valuable that an adversary with many millions of dollars worth of processing power is willing to devote it to breaking the IPsec encryption. Even in that case, there is likely an easier and much cheaper way to break into the network and achieve the same end result (social engineering, for one). Performance is the most important factor for most, and in cases when that is a concern, more care is needed when crafting a configuration.

IPsec Modes

AZTCO-FW® software supports several primary modes of IPsec operation:

Policy-based IPsec This mode uses policies to match specific combinations of traffic which are grabbed by the kernel and pushed through an IPsec tunnel. It also uses special "trap" policies to detect when traffic intends to use IPsec so that it can bring the tunnel up automatically. Only traffic specifically matching phase 2 child SA entries can use IPsec, and all traffic matching those entries will be taken over by IPsec.

This mode is the most common and is supported by nearly all third party IPsec implementations.

Route-based IPsec (VTI) Routed IPsec uses a special Virtual Tunnel Interface (VTI) for each IPsec tunnel. The VTI interface is assigned and used like other interfaces. Phase 2 entries define addresses for the tunnel interface itself, rather than policies which direct traffic to IPsec. Arbitrary traffic may cross IPsec tunnels, as traffic follows the system routing table. Static routes or dynamic routing daemons can control which traffic crosses a tunnel.

Support for routed IPsec varies by vendor.

Mobile IPsec Similar to policy-based mode, but for remote access/mobile clients.

Transport Mode This mode encrypts all traffic from the the external IP address on this firewall to the external IP address on the far side as defined in the Phase 1 settings. Since all traffic sent between the two nodes will be encrypted, other tunneling methods that do not employ encryption, such as a GIF or GRE tunnel, can be safely used by the firewall between the endpoints.

Interface Selection

In many cases, the Interface option for an IPsec tunnel will be WAN, since the tunnels are connecting to remote sites. However, there are plenty of exceptions, the most common of which are outlined in the remainder of this section.

CARP Environments

CARP type virtual IP addresses are also available in the Interface drop-down menu for use in High Availability environments (*High Availability*). In these environments, an appropriate CARP address must be chosen for the WAN where the IPsec tunnel will terminate. By using the CARP IP address, it ensures that the IPsec tunnel will be handled by the High Availability cluster member currently in MASTER state, so even if the primary firewall is down, the tunnel will connect to whichever cluster member has taken over the MASTER role.

IP Alias VIP

If multiple IP addresses are available on an interface using IP Alias type VIPs, they will also be available in this list. To use one of those IP addresses for the VPN instead, select it here.

Multi-WAN Environments

When using Multi-WAN (*Multiple WAN Connections*), pick the appropriate Interface choice for the WAN-type interface to which the tunnel will connect. If the connection will enter via WAN, pick WAN. If the tunnel will use a different WAN, choose whichever OPT WAN interface is needed. A static route will automatically be added to ensure that the traffic to the Remote Gateway routes through the appropriate WAN.

A gateway group may also be chosen from this list. A gateway group to be used with IPsec must only have *one* gateway per tier. When using a gateway group, if the first gateway goes down, the tunnel will move to the next available WAN in the group. When the first WAN comes back up, the tunnel will be rebuilt there again. If the endpoint on the far side is one that does not support multiple peer addresses, such as another firewall running AZTCO-FW software, this must be combined with a DynDNS host set using the same gateway group for failover. The DynDNS host will update the IP address as seen by the far side, so that the remote endpoint will know to accept traffic from the newly activated WAN.

Wireless Internal Protection

When configuring IPsec to add encryption to a wireless network, as described in *Additional protection for a wireless network*, choose the OPT interface which corresponds to the wireless card. When using an external wireless access point, pick the interface which is connected to the wireless access point.

Phase 1 Settings

The settings here control the phase 1 negotiation portion of the tunnel, as described previously.

General Information

Disabled Controls whether or not this tunnel (and its associated phase 2 entries) are active and used.

Key Exchange Version This can be *IKEv1*, *IKEv2*, or *Auto*. The differences are discussed in *IKE*.

IKEv1 IKEv1 is more common and widely supported, but has known issues with supporting common modern issues such as dealing with NAT or mobile clients.

IKEv2 An updated version of the protocol which has increased capabilities and security, as well as built-in support for mobile clients and NAT.

Note: IKEv2 is the best choice when both sides support it.

Auto This option uses IKEv2 when initiating, but will accept either IKEv2 or IKEv1 when responding.

Internet Protocol The protocol for the *outside* of the tunnel. That is, the protocol that will be used between the outside peer addresses. For most, this will be *IPv4* , but if both ends are capable of IPv6, that may be used instead. Whichever protocol is chosen here will be used to validate the Remote Gateway and the associated identifiers.

Note: A tunnel using IKEv1 can only carry the same protocol traffic in Phase 2 as was used for Phase 1. For example, IPv4 peer addresses restrict Phase 2 to IPv4 networks only. A tunnel using IKEv2 can carry both IPv4 and IPv6 traffic at the same time in Phase 2 no matter which protocol was used for Phase 1.

Interface This determines which part of the network will be the termination point (end point) for the IPsec tunnel. If the tunnel will be connecting to a remote server, then WAN is likely the desired setting. This can also be a virtual IP address. A gateway group can also be used for automatic failover. See *Interface Selection* earlier in this document for details on selecting the appropriate interface.

Remote Gateway The IP Address for the peer to which the tunnel will be established. This is most likely the WAN IP address of the remote firewall.

This may be set to an IP address or a fully qualified domain name. When set to use a name, the entry is periodically resolved by DNS and updated when a change is detected.

Description It is a good practice to leave notes about the purpose of a tunnel. Enter a few works to describe what this VPN tunnel is used for, or about the remote end of the tunnel. This serves as a reminder for anyone managing the firewall (present or future) as to who or what will be using the tunnel.

Authentication Method An IPsec phase 1 can be authenticated using a pre-shared key (PSK) or certificates, the Authentication Method selector chooses which of these methods will be used for authenticating the remote peer. Fields appropriate to the chosen method will be displayed on the phase 1 configuration screen.

Mutual PSK The peer is validated using a defined string. The longer the better, but since it is simple a string, there is a possibility that it can be guessed. For this reason a long/complex key is more secure when using PSK mode.

Mutual Certificate Select a CA and certificate used to verify the peer. During the phase 1 exchange, each peer sends its certificate to the other peer and then validates it against their shared CA. The CA and certificate must be created for the tunnel before attempting to setup the phase 1.

Mutual Certificate (PKCS#11) Similar to *Mutual Certificate* but the certificate is read from a locally attached PKCS#11 device.

Mutual PSK+Xauth Used with mobile IPsec and IKEv1, this selection enables xauth username and password verification along with a shared (or "group") pre-shared key.

Mutual Certificate+Xauth Used with mobile IPsec and IKEv1, this selection enables xauth username and password verification along with certificate authentication using certificates on both the client and server.

Hybrid Certificate+Xauth Used with mobile IPsec and IKEv1, this selection enables xauth username and password verification along with a certificate only on the server side. It is not quite as secure as *Mutual Certificate+Xauth* , but it is easier on the clients.

EAP-TLS Used with mobile IPsec and IKEv2, EAP-TLS verifies that certificates on the client and server are from the same shared CA, similar to *Mutual Certificate*. The client and server certificates require special handling:

- The server certificate must have the firewall hostname as it exists in DNS listed in its Common Name, and again as a Subject Alternative Name (SAN). The firewall IP address must also be listed in a SAN.
- The identifier in Phase 1 must also be set to match the firewall hostname as listed in the Common Name of the certificate.
- The client certificate must have the username listed as the common name and then again as a SAN.

The CA and server certificates must be generated before attempting to configure EAPTLS. The CA and user certificate must be imported into the client.

EAP-RADIUS Used with mobile IPsec and IKEv2, this selection performs CA verification along with username and password authentication via RADIUS. A RADIUS server must be selected on the Mobile Clients tab. Though user certificates are not necessary, EAP-RADIUS still requires that a CA and server certificate be present using the same attributes mentioned under *EAP-TLS*. The CA must be imported to the client, but no user certificate.

EAP-MSCHAPv2 Used with mobile IPsec and IKEv2, EAP-MSCHAPv2 works identically to EAP-RADIUS except the usernames and passwords are defined on the PreShared Key tab under VPN > IPsec with the Secret type set to EAP. It also requires a CA and server certificate with the same requirements listed previously. The CA must be imported to the client, but no user certificate.

Negotiation Mode (IKEv1 only) This is the type of authentication security that this tunnel will use. This can be either Main or Aggressive.

Main *Main* is the most secure mode, though it also requires more packets between the peers to accomplish a successful negotiation. It is also much more strict in its validation. This mode is best for security, but not speed.

Aggressive *Aggressive* is generally the most compatible and is the fastest mode. It is more forgiving with identifier types, and tends to be more successful when negotiating with third-party devices. It is faster because it sends all of the identifying information in a single packet, which also makes it less secure because the verification of that data is not as strict as that found in main mode.

Identifiers

My Identifier Identifies this firewall to the far side. It is best left at *My IP Address* and the firewall will fill it in as needed. In some cases an FQDN or similar may be entered so that the value is constant. So long as both sides agree on the identifier, it will work.

Peer Identifier Identifies the peer on the far side of the tunnel. It is best left at *Peer IP Address* and the firewall will fill it in as needed. In some cases an FQDN or similar may be entered so that the value is constant. So long as both sides agree on the identifier, it will work.

Identifier Types

My IP Address / Peer IP address This choice is a macro that will automatically use the IP address on the interface, or the selected VIP, as the identifier. For peers, this is the IP address from which the packets were received, which should be the Remote Gateway.

IP Address The *IP Address* option allows a different IP address to be used as the identifier. One potential use for this would be if the firewall is behind a router performing NAT. The real external IP address could be used in this field.

Distinguished Name A *Distinguished Name* is another term for a fully qualified domain name, such as host.example.com. Enter a value in that format into the box.

User Distinguished Name A *User Distinguished Name* is an e-mail address, such as vpn@example.com.

ASN.1 Distinguished Name If using *Mutual Certificate* authentication, this can be the subject of the certificate being used, or a similar string.

KeyID Tag An arbitrary string to use as the identifier.

Dynamic DNS A hostname to resolve and use as the identifier. This is mostly useful if the firewall is behind NAT and has no direct knowledge of its external IP address aside from a dynamic DNS hostname. This is not relevant or available for a

Any In cases when the remote identifier is unknown or cannot be matched, the Peer Identifier may be set to Any. This is more common on certain types of mobile configurations, but it is a much less secure choice than matching the identifier properly.

Pre-Shared Key (*Mutual PSK* authentication only) This key must be exactly the same on both VPN peers. It is case sensitive. Think of this like a "password" for the tunnel. Since this only gets entered once on each side and there is no need to remember it, it is better to make this as long and complex as possible.

Click ⟳ Generate new Pre-Shared Key to populate the field with a random long string suitable for use as a Pre-Shared Key.

> Warning: This Pre-Shared Key must be as random as possible to protect the contents of the tunnel.

My Certificate (*Mutual Certificate* authentication only) Defines the certificate which identifies this firewall. The CA which signed this certificate must be known by the peer, which may be sending them a copy of the CA certificate. If one is not shown, create or import it under System > Cert Manager on the Certificates tab.

Peer Certificate Authority (*Mutual Certificate* authentication only) Defines the CA which has signed the certificate sent by the peer. This is used to validate the peer certificate. If it does not show in the list, import it under System > Cert Manager on the Certificate Authorities tab.

Phase 1 Encryption algorithms

There are many options for encryption algorithms on both phase 1 and phase 2.

Encryption choices depend on the device to which the tunnel will connect, and the hardware available in this firewall. Generally speaking, AES-GCM is the most desirable cipher. When connecting to third party devices, 3DES (also called "Triple DES") is a common choice as it may be the only option the other end supports, though it is considered weak and should be avoided when possible.

Phase 1 Encryption Options Multiple combinations of these options can be defined using the ✚ Add Algorithm button to add another line.

Encryption Algorithm If both sides support AES-GCM, use *AES128-GCM* with a *128* bit Key Length. This will combine strong encryption and hashing together and can be accelerated by AES-NI. Failing that, use *AES* With a Key Length of *128*. If the peer does not support any of these, use the strongest available option supported by the peer.

Hash Algorithm Hash algorithms are used with IPsec to verify the authenticity of packet data and as a Pseudo-Random Function (PRF). These hash algorithms may also be referred to with HMAC (Hash Message Authentication Code) in the name in some contexts, but that usage varies depending on the hardware or software in use.

When using AES-GCM, this is used solely as a PRF because AES-GCM already performs hashing internally. The best choice for use with AES-GCM is *AES-XCBC*. If a different type of Encryption

Algorithm is in use, then use *SHA256* if possible. If the peer does not support any of these, use the strongest available option supported by the peer.

PRF Specifically set a manual PRF different than the one the IPsec daemon would choose automatically based on the Hash Algorithm. This control is hidden by the GUI unless PRF Selection is enabled in the Advanced Options section at the bottom of the page.

DH Key Group The best practice is to use DH Group *14 (2048 bit)* or higher if both sides support it. Avoid using groups 1, 2, 22, 23, and 24 as they do not provide sufficient security. As with the other options, if the suggested value is not supported by the peer, use the strongest available option.

Expiration and Replacement

The total lifetime for Phase 1 defines how often the connection will be rekeyed or reauthenticated by the IPsec daemon, in seconds.

In most cases, a tunnel should use either Rekey Time or Reauth Time, plus Over Time. The specific values depend on the IKE mode and which mechanisms are supported by both endpoints.

28800 total seconds is a good balance of frequent rekeying without being too aggressive. Place approximately 90% of this total in either Rekey Time or Reauth Time, and the remaining amount in Over Time.

Rekey Time Time, in seconds, before the IPsec daemon attempts attempts to establish a new set of keys for the IKE SA. Only supported by IKEv2, and is the best choice for use with IKEv2.

Rekey works without interruption, allowing both endpoints to seamlessly change to new keys on the fly. This is optimal, but implementation quality varies by vendor.

Leave blank or enter a value of 0 to disable rekeying.

Normally both sides will rekey as needed, but if the tunnel often fails when a rekey event occurs, try disabling this feature on one side.

Note: Some clients, especially Windows clients behind NAT, misbehave when they receive a rekey request. In those cases it is safer to allow the client to initiate the rekey by disabling the option on the server.

Reauth Time Time, in seconds, before an IKE SA is torn down and recreated from scratch by the IPsec daemon, including authentication. Supported by IKEv1 and IKEv2, but should be avoided with IKEv2 where possible.

This process can be disruptive to traffic flow unless all peers support IKEv2 make-before-break (*Advanced IPsec Settings*) and overlapping IKE SA entries.

Leave blank or enter a value of 0 to disable reauthentication.

Over Time Hard IKE SA life time, in seconds, after which the IKE SA will expire. This time is relative to reauthentication and rekey time.

If empty, defaults to 10% of whichever timer is higher (reauth or rekey).

Advanced Options

Responder Only When set, then IPsec daemon will not attempt to initiate the tunnel. The tunnel will only be established by an initiation attempt from the far side. Also, if DPD detects that the tunnel has failed, the tunnel will be left down rather than restarted, leaving it up to the far side to reconnect.

Child SA Close Action Controls how the IPsec daemon behaves when a child SA (P2) is unexpectedly closed by the peer.

Default Retains the default behavior based on other settings for the tunnel.

Close connection and clear SA Removes the child SA and does not attempt to establish a new SA. This is the desired behavior when acting in a Responder Only or mobile IPsec role.

Restart/Reconnect Immediately attempts to reconnect the child SA. This ensures that the tunnel reestablishes properly in cases that do not support trap policies, such as routed IPsec (VTI). Set this on one side only if the tunnel does not reconnect after it disconnects, rekeys, or reauthenticates.

> Warning: This option must not be set on both peers! Otherwise both peers will attempt to initiate and hold open multiple copies of each child SA.

Close connection and reconnect on demand Clears the child SA and reinstalls trap policies to watch for interesting traffic. Will reestablish the tunnel on demand when traffic attempts to cross the tunnel.

This option is not compatible with modes which do not support trap policies, such as routed IPsec (VTI).

NAT Traversal (IKEv1 Only) Also known as NAT-T. NAT Traversal encapsulates ESP traffic for IPsec inside of UDP packets, to more easily function in the presence of NAT. If this firewall or the firewall on the other end of the tunnel is behind a NAT device, then NAT Traversal will likely be necessary for the tunnel to function properly.

Auto (Default) Allows the IPsec daemon to detect and use NAT Traversal automatically when it determines one or both peers is behind NAT.

Force Instructs the IPsec daemon to always use NAT Traversal for the tunnel. This can help if there is a known issue detecting NAT, or with issues carrying ESP traffic between the two endpoints even when neither side is behind NAT.

IKEv2 integrates NAT Traversal natively so the option is unnecessary in that case.

MOBIKE An extension to IKEv2 which handles multi-homed clients and clients which roam between different IP addresses. This is primarily used with mobile clients to allow them to switch remote addresses while keeping the connection active. Leave disabled unless the remote peer must change addresses dynamically.

Gateway Duplicates When set, the GUI validation allows multiple Phase 1 configurations to the same remote endpoint.

This option also disables automatic static routes to the peer via specific WAN gateways. Traffic will follow the default route, not the tunnel interface, unless manual static routes redirect the traffic.

Split Connections (IKEv2 Only) When an IKEv2 tunnel has multiple Phase 2 definitions, by default the settings are collapsed in the IPsec configuration such that all P2 combinations are held in a single child SA.

Split Connections changes this behavior to be more like IKEv1 where each P2 is its configured by the daemon as own separate child SA.

Certain scenarios require this behavior, such as:

- The remote peer does not properly handle multiple addresses in single traffic selectors. This is especially common in Cisco equipment.
- Each child SA must have unique traffic selector or proposal settings. This could be due to the peer only allowing specific combinations of local/remote subnet pairs or different encryption options for each child SA.

PRF Selection When set, the GUI enables a control to specifically set a Pseudo-Random Function (PRF) rather than allow the IPsec daemon to choose one automatically based on the selected Hash Algorithm. Can be useful in combination with AEAD encryption algorithms such as AES-GCM.

Dead Peer Detection Dead Peer Detection (DPD) is a periodic check that the host on the other end of the IPsec tunnel is still alive. This detects when an IPsec peer has lost connectivity or otherwise is unreachable. If a DPD check fails, the tunnel is torn down by removing its associated SAD entries and a fresh negotiation is attempted.

The default settings are sufficient for most connections. Increase the values for bad quality links to avoid tearing down a usable, but lossy, tunnel.

Delay Time between DPD probe attempts. The default of 10 is best.

Max Failures Number of failures before the peer is considered down. The default of 5 is best.

Note: The default values of 10 seconds and 5 failures will result in the tunnel being considered down after approximately one minute.

Phase 2 Settings

The phase 2 settings for an IPsec tunnel govern how the tunnel handles traffic (e.g. policy-based or route-based, see *IPsec Modes*) as well as the encryption of that traffic.

Phase 2 entries are used in a few different ways, depending on the IPsec configurations:

- For policy-based IPsec tunnels, this controls which subnets will enter IPsec. Multiple phase 2 definitions can be added for each phase 1 to allow using multiple subnets inside of a single tunnel.
- For route-based IPsec, this controls the VTI interface addresses.
- For mobile IPsec this primarily controls the encryption for phase 2, but can also optionally be used by the IPsec daemon or export utilities to generate a list of networks to the clients for use in split tunneling.

Each Phase 2 entry has the following options:

Disabled An on/off switch for this Phase 2 entry only.

Mode The IPsec Mode for this Phase 2 entry, which controls how the tunnel handles traffic. See *IPsec Modes* for more detailed explanations of each type of mode.

With policy-based IKEv1 tunnels, this must match the outer protocol of the tunnel, for example an IPv4 peer would be Tunnel IPv4. Policy-based IKEv2 tunnels can have either/or (or both).

Tunnel IPv4 A policy-based tunnel that will carry traffic between IPv4 networks matching the specified Local Network and Remote Network.

Tunnel IPv6 A policy-based tunnel that will carry traffic between IPv6 networks matching the specified Local Network and Remote Network.

Transport Encrypts all traffic between the endpoints. The Local Network and Remote Network are not set for transport mode, it assumes the addresses based on the phase 1 settings.

Routed (VTI) Routed IPsec using Virtual Tunnel Interfaces. The Local Network and Remote Network define the addresses used by the firewall for the VTI interface. Typically only one Phase 2 entry is present for each address family (e.g. one for IPv4, one for IPv6)

See *Routed IPsec (VTI)* for more information.

Local Network

Tunnel Mode Defines which subnet or host can be accessed from the other side of the VPN tunnel. This is typically the LAN or other internal subnet for the VPN, but can also be a single IP address if only one client needs to use the tunnel. The Type selector is pre-loaded with choices for each interface (e.g. *LAN subnet*), as well as *Address* and *Network* choices that allow entering an IP address or subnet manually.

Most often this is set to *LAN subnet*, meaning the entire LAN will be accessible from the remote network.

NAT/BINAT Sets a *different* subnet or address which is used by IPsec to perform NAT on the local network addresses to make them appear to the remote peer as a different subnet.

Set to *None* to disable NAT for the tunnel.

For more details, see *NAT with IPsec Phase 2 Networks*.

Routed (VTI) Mode Sets the local IP address and subnet mask of the ipsecX interface. Remote

Network

Tunnel Mode (Non-mobile only) Specifies the IP Address or Network which exists on the other (remote) side of the VPN. This field operates similarly to the Local Network option.

Routed (VTI) Mode Sets the remote IP address for the ipsecX interface tunnel network (the remote address of the VTI).

Description A description for this Phase 2 entry. Shows up in the IPsec status for reference.

Phase 2 Proposal (Child SA)

Protocol Controls how IPsec protects its traffic.

ESP (Encapsulating Security Payload) Encrypts traffic before sending it to the peer.

In nearly all circumstances, ESP is the correct choice.

AH (Authenticated Header) Provides assurance the traffic came from a trusted source but does not provide encryption. Rarely used in practice.

Note: With automatic VPN rules (*Disable Auto-added VPN rules*), the firewall automatically passes the appropriate ESP or AH protocol traffic from the remote endpoint. If automatic VPN rules are disabled, add manual rules to pass the traffic instead.

Encryption algorithms Sets the encryption algorithms used when negotiating Phase 2 child SA entries with peers. Must match values available to and configured on the peer.

In systems with AES-NI, the fastest and most secure choice is AES-GCM, if it is supported by the peer. If AES-CGM is used, do not select any options for Hash Algorithms in Phase 2.

If AES-NI cannot be used both both peers, use AES with a 128-bit or higher key length.

This set of controls allows for multiple selections so that multiple choices will be accepted when acting as a responder, or proposed when working as an initiator. The best practice is to only select a single desired cipher on both peers, but in some cases, such as mobile clients, selecting multiple will allow a tunnel to work better in both a responder and initiator role.

Hash algorithms Controls which hash algorithms are used when negotiating phase 2 child SA entries with peers. Must match values available to and configured on the peer.

As with the Encryption Algorithms, multiple hashes may be selected. The best practice is to select a single desired choice if possible. For more discussions on the quality of the various hash types, see *Phase 1 Settings*.

The optimal choice for speed and security is SHA256, unless using AES-GCM for the Encryption Algorithm. With AES-GCM, no Hash Algorithm should be selected as AES-GCM performs hashing on its own.

PFS key group Perfect Forward Secrecy (PFS) provides keying material with greater entropy, hence improving the cryptographic security of the connection, at the cost of higher CPU usage during rekeying.

The options have the same properties as the DH key group option in phase 1 (See *DH key group*), and some products also refer to them as “DH” values even in Phase 2.

The optimal choice for speed and security is *14 (2048 bit)*. The default is *off*.

Lifetime The lifetime for which the negotiated keys will be valid. One hour (3600) is a good setting. Do not set this to too high (e.g. more than about a day: 86400) as doing so will give people more time to crack the key. Don’t be over paranoid either; there is no need to set this to 20 minutes either.

Automatically ping host For use on non-mobile tunnels, this option tells the firewall to initiate a ping periodically to the specified IP address. This option only works if the firewall has an IP address inside of the Local Network for this Phase 2 entry and the value of the ping host here must be inside of the Remote Network.

See *Configuring IPsec Keep Alive* for additional information.

14.7.2 IPsec Tunnel List

The IPsec page located at VPN > IPsec allows management of IPsec VPN tunnels. A brief summary of existing tunnel settings is also displayed on this page.

Each IPsec tunnel will have one phase 1 definition, and one or more phase 2 definitions.

Phase 1 definitions handle how the tunnel connects to the remote peer. This includes the remote gateway, authentication information such as identifiers and pre-shared keys or certificates, NAT Traversal and DPD settings.

After adding a phase 1 definition, click the larger + button underneath a phase 1 entry to display and manage its phase 2 entries.

Phase 2 definitions handle how local/internal networks are sent across a tunnel. Multiple local subnets (or individual hosts) can be used on a single IPsec tunnel by adding multiple Phase 2 entries.

Phase 2 definition settings include the local and remote networks for traffic which will traverse the tunnel, and phase 2 encryption proposal settings.

See also:

IPsec - All other IPsec articles.

14.7.3 NAT with IPsec Phase 2 Networks

AZTCO-FW® software supports for NAT on policy-based IPsec Phase 2 entries to make the local network appear to the remote peer as a different subnet or address. This can be used to work around subnet conflicts or connect to vendors without renumbering a local network.

Warning: NAT is not currently compatible with route-based VTI IPsec tunnels.

Configuration

NAT is configured by the NAT/BINAT Translation options on an IPsec Phase 2 entry in tunnel mode, in combination with the Local Network settings.

Local Network Values of Type and Address specify the actual local network (e.g. LAN subnet).

NAT/BINAT Translation Values of Type and Address specify the translated network visible to the far side.

NAT Types

There are two main modes for NAT with IPsec:

Binat - 1:1 NAT When both the actual and translated local networks use the same subnet mask, the firewall will directly translate the networks to one another inbound and outbound. Can also be used for single addresses.

This allows remote host to directly contact local hosts using their equivalent NAT addresses, provided that IPsec rules allow the traffic to pass.

NAT - Overload/PAT Style If the Local Network is a subnet, but the NAT/BINAT Translation address is set to a single IP address, then a 1:many NAT (PAT) translation is set up that works like an outbound

NAT rule on WAN. All outbound traffic will be translated from the local network to the single IP address in the NAT field.

Note: Inbound traffic from the remote network to individual local hosts is not possible in this mode.

Warning: NAT+IPsec cannot be configured between two different sized subnets (e.g. It cannot NAT a /24 subnet to a /27 subnet).

Example

Consider an IPsec tunnel to a Vendor which requires 172.16.5.0/24 for the network on this firewall. However, the LAN is actually 192.168.1.0/24, and renumbering is not feasible.

To accommodate this scenario, set the Phase 2 values as follows:

Local Network

Type *Network*

Address 192.168.1.0/24

NAT/BINAT Translation

Type *Network*

Address 172.16.5.0/24

Firewall Rules

NAT is processed before firewall rules, so firewall rules on the IPsec tab refer to the network in Local Network.

Remote End Notes

The far side of the tunnel does not need any knowledge of the actual Local Network. Their tunnel is built between their local network and the NAT/BINAT Translation value.

Packet Capturing Quirk

In a packet capture, the Local Network addresses are shown on outbound traffic, not the translated address. This does not indicate any problem.

14.7.4 Routed IPsec (VTI)

Route-based IPsec is an alternative method of managing IPsec traffic. It uses if_ipsec(4) from FreeBSD 11.1+ for Virtual Tunnel Interfaces (VTI) and traffic is directed using the operating system routing table. It does not rely on strict kernel security association matching like policy-based (Tunneled) IPsec.

A routed IPsec tunnel creates an ipsecXXXX interface at the operating system level and this interface has its own IP address. The ipsecXXXX interface must be assigned so it can be used for purposes such as static or dynamic routing, daemon binding, traffic monitoring, and so on.

Once assigned, the IPsec interface also gains an automatic gateway which provides policy routing and gateway group capabilities.

Note: Routed IPsec is not replacing traditional tunneled IPsec, both may be used. The choice is up to the user when creating an IPsec Phase 2 entry.

Note: Routed IPsec is available on firewalls running AZTCO-FW® software version 2.4.4-RELEASE or later.

Prerequisites

First, pick a transit network. This is similar to choosing a tunnel network for an OpenVPN instance. Typically this is a /30 network in an unused subnet. This example uses 10.6.106.0/30.

IPsec Configuration

- Create an IPsec Phase 1 entry as usual.
- Create a Phase 2 entry under this Phase 1, set with. . .
  - Set Mode to *Routed (VTI)*
  - Enter 10.6.106.1 for the Local Network Address
  - Enter 10.6.106.2 for the Remote Network Address
  - Add a useful Description
  - Set the Proposal settings as needed
- Click Save, then click Apply Changes

IPsec Interface Assignment

- Navigate to Interfaces > Assignments
- Pick the new ipsecX interface from the Available Network Ports list
- Click + Add
- Note the new interface name, e.g. *OPT1*
- Navigate to Interfaces > [New Interface Name]
- Check Enable
- Give the interface a more suitable name using the Description field (e.g. VTI_FOO)
- Leave the IPv4 Configuration Type and IPv6 Configuration Type set to *None*

- Click Save, then click Apply Changes

A gateway is created automatically and can be used for static routing, policy routing, and so on.

At this point the interface is available for use like any other interface. It can be used for packet captures, traffic graphs, binding daemons, routing protocols, and other tasks never before possible with IPsec on AZTCO-FW software!

Routing

Until routing is configured, no traffic will attempt to cross the IPsec tunnel except for gateway monitoring probes, if they are enabled.

Static Routes

To setup static routes, navigate to System > Routing, Static Routes tab. Add a new route there using the assigned IPsec interface gateway.

Policy Routes

To policy route traffic across a routed IPsec tunnel, use the assigned IPsec interface gateway in firewall rules as usual for policy routing.

See also:

Policy Routing Configuration

Dynamic Routes

The assigned IPsec interface can be used in dynamic routing daemons such as FRR, Quagga, and OpenBGPD. BGP and OSPF can both operate across routed IPsec interfaces.

Routed IPsec Firewall Rules

Routed IPsec traffic appears to the OS on both the specific IPsec interface and the enc0 interface, which is governed by the rules on the IPsec tab. Though a tab appears for the assigned interface, traffic must be passed on the IPsec tab.

Caveats

Routed IPsec works best when both sides support routed IPsec. It can still work when only one side supports routed IPsec, but most of its benefits are lost.

Rather than managing IPsec Phase 2 entries, routes must be managed instead. Since this can be automated with dynamic routing protocols this is not a large concern.

Firewall rule processing can be confusing, as mentioned in *Routed IPsec Firewall Rules*. This is still undergoing testing, but likely means that reply-to will not function. There are also known issues with NAT, notably that NAT to the interface address works but 1:1 NAT or NAT to an alternate address does not work.

14.7.5 IPsec and firewall rules

When an IPsec tunnel is configured, AZTCO-FW® automatically adds hidden firewall rules to allow UDP ports 500 and 4500, and the ESP protocol from the Remote gateway IP address destined to the Interface IP address specified in the tunnel configuration. When mobile client support is enabled the same firewall rules are added except with the source set to *any*. To override the automatic addition of these rules, check Disable all auto-added VPN rules under System > Advanced on the Firewall/NAT tab. When that box is checked, firewall rules must be manually added for UDP 500, UDP 4500, and ESP to the appropriate WAN interface.

Traffic initiated from the remote end of an IPsec connection is filtered with the rules configured under Firewall > Rules on the IPsec tab. Here restrictions may be placed on resources made accessible to remote IPsec users. To control what traffic can be passed from local networks to the remote IPsec VPN connected devices or networks, the rules on the local interface where the host resides control the traffic (e.g. connectivity from hosts on LAN are controlled with LAN rules).

14.7.6 Using IPsec with Multiple Subnets

On current versions of AZTCO-FW® software, additional subnets are handled by adding an additional Phase 2 entry to cover the path to pass through the tunnel.

For example, for 172.16.0.0/24 and 172.16.1.0/24 at Site A, and 10.0.0.0/24 at Site B, define two Phase 2 entries on both sides:

On the Site A Firewall:

```
172.16.0.0/24 to 10.0.0.0/24
172.16.1.0/24 to 10.0.0.0/24
```

On the Site B Firewall:

```
10.0.0.0/24 to 172.16.0.0/24
10.0.0.0/24 to 172.16.1.0/24
```

This works for any additional networks on either side (VPN subnets, networks on the other end of VPNs connected to the remote router, etc).

If the equipment to which the tunnel connects does not support multiple Phase 2's, it may be necessary to employ supernetting/CIDR summarization (See below) to fit the networks into a single Phase 2.

Supernetting Example

At Site A, there is one subnet, 10.0.0.0/24. This should reach 192.168.0.0/24, 192.168.1.0/24, and 192.168.2.0/24 at Site B.

Due to the "closeness" of the subnets, they could be grouped into a larger network in the tunnel definition: 192.168.0.0/22 (This would also include 192.168.3.0/24)

14.7.7 Advanced IPsec Settings

The Advanced Settings tab under VPN > IPsec contains options to control, in general, how the IPsec daemon behaves and how traffic is handled with IPsec.

IPsec Logging Controls These options control which areas of the IPsec daemon generate log messages and their level of detail. For information on viewing the log, see *IPsec Logs*.

In most cases the optimal settings are the default: IKE SA, IKE Child SA, and Configuration Backend set to *Diag*, and all others set to *Control*.

Configure Unique IDs as Controls how the IPsec daemon treats new connections with an identifier which matches an existing connection. In most cases a new connection is intended to replace an older connection, but certain use cases such as mobile clients may require multiple connections from the same remote identifier.

Yes (Replace) The new connection is accepted by the IPsec daemon and it replaces the old connection, which is disconnected.

No The new connection is accepted and the old connection is replaced only if the peer sends an INITIAL_CONTACT notification.

Never The new connection is always allowed, and INITIAL_CONTACT notifications are ignored.

Keep The new connection is rejected and the old connection remains active.

IP Compression Propose support for IPComp compression.

> Warning: Though the option is present in the GUI, the underlying operating system does not yet fully support IP compression.

Strict Interface Binding When set, the IPsec daemon configuration binds only to the interfaces required by the configuration, rather than binding to all interfaces.

This option is more secure but is known to break with interfaces which have dynamic IP addresses. Only enable this option in environments where it has been lab tested and proven to work as intended.

Unencrypted Payloads in IKEv1 Main Mode Some IPsec implementations send the third Main Mode message unencrypted, probably to find the PSKs for the specified ID for authentication. This is similar to Aggressive Mode, and has the same security implications: A passive attacker can sniff the negotiated Identity, and start brute forcing the PSK using the HASH payload. The best practice is to keep this option disabled unless the implications are fully understood and compatibility to such devices is required (for example, some SonicWall devices).

MSS Clamping Enable maximum segment size clamping on TCP flows over IPsec tunnels. This helps overcome problems with path MTU discovery (PMTUD) on IPsec VPN links.

This is useful is large TCP packets have problems traversing the VPN, or if slow/choppy connections across the VPN are observed by users. Ideally it should be set to the same value on both sides of the VPN, but traffic will have MSS clamping applied in both directions.

Enable When set, the Maximum MSS option is available and its value is used by the firewall configuration.

Maximum MSS The maximum segment size set in TCP packets flowing across IPsec VPN tunnels. Defaults to 1400. Must be low enough to account for the overhead of IPsec and the MTU of the link, but no so low that unnecessarily small segments are sent as that can be inefficient.

Enable Cisco Extensions Enables the Unity plugin which provides support for Cisco Extensions such as Split-Include, Split-Exclude, and Split-DNS for IKEv1 XAuth mobile clients. This allows clients which support these extensions to obtain values automatically when connecting to a mobile IPsec VPN.

Strict CRL Checking When set, the IPsec daemon requires availability of a fresh CRL for peer authentication based on certificate signatures to succeed. Primarily useful when the CRL is obtained dynamically (e.g. OCSP).

Make Before Break Controls whether IKEv2 Reauthentication uses Make-before-Break or Breakbefore-Make when an IKE Security Association (SA) expires. Must be supported by both peers.

Only relevant for IKEv2 tunnels using reauthentication, it does not affect IKEv1 tunnels or IKEv2 tunnels set to rekey.

Break-before-Make (Unchecked, Default) Deletes IKE and Child SAs before reauthenticating and making a new set of SAs. This behavior is standard and well-supported, but disruptive as there is a small gap between the old and new SA set in which IPsec connectivity is unavailable.

Make-before-Break (Checked) Reauthenticates and makes a new SA set before deleting the old SA set. This eliminates the connectivity disruption, but requires that both endpoints support overlapping IKE and Child SA entries.

Asynchronous Cryptography Allows cryptographic framework jobs to be dispatched in a multithreaded manner to increase performance. Jobs are handled in the order they are received so that packets will be reinjected in the correct order.

> Warning: This option can increase performance, but may be unstable on certain hardware. When enabling this option, test connectivity during a maintenance window to ensure proper behavior. See Bug #8964 for details.

Auto-exclude LAN Address Set up an automatic IPsec bypass for traffic to and from the LAN subnet, so it does not get captured by policy-based IPsec.

Additional IPsec Bypass Configures additional manual IPsec bypass behavior. When set, the GUI exposes the IPsec Bypass Rules control.

IPsec Bypass Rules Custom rules which allow traffic matching combinations of Source Address and Destination Address pairs to be excluded from IPsec policies.

Source Address The source address or network to exclude, and its mask.

Destination Address The corresponding destination address or network to exclude, and its mask.

> Note: These values are considered *together*. A packet must match both the source and destination to bypass IPsec policies.

These rules are useful to exclude traffic between multiple local networks, especially when a policy-based IPsec tunnel is set to use 0.0.0.0/0 as the remote network.

14.7.8 Configuring IPsec Keep Alive

Any IP address within the Remote Network of this phase 2 definition may be used. It does not have to reply or even exist, simply triggering traffic destined to that network periodically will keep the IPsec connection up and running.

For this feature to work, the firewall must have an IP address assigned inside the Local Network. Otherwise it cannot generate the necessary traffic to bring up the tunnel.

14.7.9 Mobile IPsec

Choosing a Mobile IPsec Style

Currently only one type of mobile IPsec may be configured at a time, though there are multiple different styles to choose from.

- *IKEv2 with EAP-TLS* for per-user certificate authentication
- *IKEv2 with EAP-MSCHAPv2* for local username and password authentication
- *IKEv2 with EAP-RADIUS* for remote username and password authentication
- *Xauth+PSK* for local or remote username and password authentication
- Xauth+RSA for certificates and local or remote username and password authentication
- *Pre-Shared Key* for basic IPsec connectivity from older clients
- *L2TP/IPsec* for local or remote username and password authentication with clients that do not support one of the above methods.

See also:

- *Configuring IPsec IKEv2 Remote Access VPN Clients on Windows*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on Ubuntu*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on Android*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on OS X*
- *Configuring IPsec IKEv2 Remote Access VPN Clients on iOS*

As of this writing, most current operating systems natively support IKEv2 or can use an app/add-on. It is currently the best choice, and will be the one demonstrated later in this chapter. Windows 7 and later, MAC OS X 10.11 (El Capitan) and later, iOS 9 and later, and most Linux distributions have support built in for IKEv2. There is a simpleto-use strongSwan IKEv2 app for Android 4.x and later.

Note: All IKEv2 types require a certificate structure including at least a Certificate Authority and a Server Certificate, and in some cases user certificates. For more information on Certificates, see *Certificate Management*. Clients can be very picky about certificate attributes, so pay close attention to this chapter when creating the certificate structure.

Warning: When generating a Server Certificate for use with IKEv2, the Common Name of the certificate must be the firewall's name as it exists in DNS. The name must be repeated again as an FQDN type Subject Alternative Name (SAN). The IP address of the firewall must also be present as an IP Address type SAN. This information will be repeated later in the chapter, but requires extra emphasis due to its importance. See *Create a Server Certificate*

IKEv2 with EAP-MSCHAPv2

With support for IKEv2 now widespread, it is an ideal choice for current operating systems. Though there are several variations, EAP-MSCHAPv2 is the easiest to configure since it does not require generating or installing per-user certificates and does not require a working RADIUS server. The CA Certificate must still be installed onto the client as a trusted root certificate.

EAP-MSCHAPv2 allows for username and password authentication using passwords stored on the Pre-Shared Keys tab under VPN > IPsec. These passwords are stored in plain text, so it is not as secure as using a RADIUS server, though it is more convenient.

See also:

- *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2*

IKEv2 with EAP-RADIUS

EAP-RADIUS works identically to EAP-MSCHAPv2 except that user authentication happens via RADIUS. When EAP-RADIUS is chosen, a RADIUS server must on the Mobile Clients tab. The RADIUS server must accept and understand EAP requests and it must also allow MSCHAPv2. Password security is left up to the RADIUS server.

EAP-RADIUS is typically the best choice when a RADIUS server is available.

See also:

- *IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS*

IKEv2 with EAP-TLS

EAP-TLS uses per-user certificate authentication instead of username and password authentication. As such, EAPTLS requires generating certificates for each user, which makes it a bit more cumbersome from an administration standpoint. Certificates are validated against the CA similar to OpenVPN. The CA certificate, user certificate and its associated key must all be imported to the client properly.

Warning: When creating user certificates, the username must be used as the certificate common name and again as a DNS/FQDN type Subject Alternative Name. If the same name is not present in both places, clients may not be validated properly.

See also:

- *IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS*

IKEv1 with Xauth and Pre-Shared Keys

Xauth+PSK works on a majority of platforms, the notable exception being current versions of Android. Windows XP through Windows 8 can use the Shrew Soft client, but Windows 10 does not currently work with any client. OS X and iOS can use their built-in client to connect.

Note: When using Xauth, local users must exist in the User Manager and those users must have the *User - VPN IPsec Xauth Dialin* privilege.

See also:

- *IPsec Remote Access VPN Example Using IKEv1 with Xauth*

IKEv1 with Xauth and RSA Certificates

Xauth+RSA works in most of the same conditions as Xauth+PSK, though it does in fact work on current Android devices. Certificates must be made for each user, and the certificates must be imported into the clients.

IKEv1 with Pre-Shared Keys Only

Pre-Shared Key only IPsec VPNs for mobile IPsec have become rare in modern times. Support was not very common, only found in the Shrew Soft client, some very specific Android versions (such as those from Motorola), and in other third-party clients. They are not very secure, and are no longer recommended for general use. The only time they may be needed is in cases when the far side cannot support any other method.

See also:

- *IPsec Remote Access VPN Example Using IKEv1 with Pre-Shared Keys*

L2TP/IPsec (IKEv1)

L2TP/IPsec is a unique combination that, unfortunately, does not work very well in most cases. In this style of setup, Mobile IPsec is setup to accept Transport Mode connections which secure all traffic between the public IP address endpoints. Across this transport channel, an L2TP connection is made to tunnel user traffic in a more flexible way. Though support for this model is found in most versions of Windows, MAC, Android, and other Operating Systems, they are all picky in different incompatible ways about what will work.

For example, the Windows client does not work properly when the client system is behind NAT, which is the most common place that a VPN client would find itself. The problem is in an interaction between the client and the IPsec daemon used on AZTCO-FW®, strongSwan. The strongSwan project states that it is a bug in the Windows client, but it is unlikely to be fixed since both strongSwan and Windows have focused their mobile client efforts on more modern and secure implementations such as IKEv2 instead.

Warning: L2TP/IPsec should be avoided when possible.

See also:

- *L2TP/IPsec Remote Access VPN Configuration Example*

See also:

- *Remote Access Mobile VPN Client Compatibility*

Mobile IPsec allows creation of a so-called "Road Warrior" style VPN, named after the variable nature of anyone who is not in the office that needs to connect back to the main network. It can be a sales person using Wi-Fi on a business trip, the boss from his limo via 3G modem, or a programmer working from their broadband line at home. Most of these will be forced to deal with dynamic IP addresses, and often will not even know the IP address they have. Without a router or firewall supporting IPsec, a traditional IPsec tunnel will not work. In telecommuting scenarios, it's usually undesirable and unnecessary to connect the user's entire home network to the office network, and doing so can introduce routing complications. This is where IPsec Mobile Clients are most useful.

There is only one definition for Mobile IPsec on AZTCO-FW®, so Instead of relying on a fixed address for the remote end of the tunnel, Mobile IPsec uses some form of authentication to allow a username to be distinguished. This could be a username and password with IKEv2 and EAP or xauth, or a per-user Identifier and Pre-Shared Key pair, or a certificate.

14.7.10 Testing IPsec Connectivity

The easiest test for an IPsec tunnel is a ping from one client station behind the firewall to another on the opposite side. If that works, the tunnel is up and working properly.

As mentioned in AZTCO-FW*-initiated Traffic and IPsec*, traffic initiated from the AZTCO-FW® firewall will not normally traverse the tunnel without extra routing, but there is a quick way to test the connection from the firewall itself by specifying a source when issuing a ping.

There are two methods for performing this test: the GUI, and the shell.

Specifying a Ping Source in the GUI

In the GUI, a ping may be sent with a specific source as follows:

- Navigate to Diagnostics > Ping
- Enter an IP address on the remote router within the remote subnet listed for the tunnel in the Host field (e.g. 10.5.0.1)
- Select the appropriate IP Protocol, likely *IPv4*
- Select a Source Address which is an interface or IP address on the local firewall which is inside the local Phase 2 network (e.g. Select *LAN* for the LAN IP address)
- Set an appropriate Count, such as the default *3*
- Click Ping

If the tunnel is working properly, ping replies will be received by the firewall from the LAN address at Site B. If replies are not received, move on to the *Troubleshooting IPsec VPNs* section.

If the tunnel was not established initially, it is common for a few pings to be lost during tunnel negotiation, so choosing a higher count or re-running the test is a good practice if the first attempt fails.

Specifying a Ping Source in the Shell

Using the shell on the console or via ssh, the ping command can be run manually and a source address may be specified with the *-S* parameter. Without using - *S* or a static route, the packets generated by ping will not attempt to traverse the tunnel. This is the syntax for a proper test:

```
# ping -S <Local LAN IP> <Remote LAN IP>
```

Where the *Local LAN IP* is an IP address on an internal interface within in the local subnet definition for the tunnel, and the *Remote LAN IP* is an IP address on the remote router within the remote subnet listed for the tunnel. In most cases this is simply the LAN IP address of the respective AZTCO-FW firewalls. Given the site-to-site example above, this is what would be typed to test from the console of the Site A firewall:

```
# ping -S 10.3.0.1 10.5.0.1
```

If the tunnel is working properly, ping replies will be received by the firewall from the LAN address at Site B. If replies are not received, move on to the *Troubleshooting IPsec VPNs* section.

14.7.12 Accessing Firewall Services over IPsec VPNs

With an out of the box configuration, it is not possible to query SNMP on the LAN interface of a remote AZTCO-FW® instance over an IPsec VPN connection.

Fred Wright explained in a post to the m0n0wall mailing list on September 12, 2004 why this is, and it's the same reason in AZTCO-FW software.

> Due to the way IPsec tunnels are kludged into the FreeBSD kernel, any traffic *initiated* by m0n0wall to go through an IPsec tunnel gets the wrong source IP (and typically doesn't go through the tunnel at all as a result). Theoretically this *shouldn't* be an issue for the *server* side of SNMP, but perhaps the server has a bug (well, deficiency, at least) where it doesn't send the response out through a socket bound to the request packet. You can fake it out by adding a bogus static route to the remote end of the tunnel via the m0n0wall's LAN IP (assuming that's within the near-end tunnel range). A good test is to see whether you can ping something at the remote end of the tunnel (e.g. the SNMP remote) *from* the m0n0wall. There's an annoying but mostly harmless side-effect to this - every LAN packet to the tunnel elicits a no-change ICMP Redirect.

Most notably this is a problem for UDP services. UDP services reply using the "closest" address to the client as seen from the perspective of the system routing table. Without a route present, that ends up being the IP address of the default gateway on WAN.

To add this route in the AZTCO-FW webGUI, perform the following configuration:

- Navigate to System > Routing on the Gateways tab
- Click + to add a gateway
- Select *LAN* for the Interface
- Enter the LAN IP address in the Gateway field
- Check Disable Gateway Monitoring
- Click Save
- Click Apply Changes
- Navigate to the Static Routes tab
- Click +
- Enter the remote VPN network in the Destination Network box
- Select the LAN IP Gateway that was created before
- Add a Description if desired
- Click Save

- Click Apply Changes

To perform a quick test with ping from the console or ssh, adjust the ping source to enable traffic to traverse the tunnel like so:

```
ping -S <AZTCO-FW LAN ip> <remote IP address>
```

If the AZTCO-FW LAN address is 192.168.1.1, and that IP is a part of the subnet defined for the IPsec tunnel, to ping 10.0.0.1 on the other side, do this:

```
ping -S 192.168.1.1 10.0.0.1
```

Another alternative, depending on the version, would be to change the interface binding of the target service so that it only listens on the LAN IP address (or the IP address of the internal network on the local end of the VPN) on the firewall. The interface binding for SNMP, NTP, the DNS Forwarder, and several other services can be set in this way.

14.7.13 IPsec and IPv6

IPsec is capable of connecting to a tunnel over IPv4 or IPv6 phase 1 peer addresses, but with IKEv1 the tunnel can only contain the same type of traffic inside the tunnel phase 2 definition that is used to pass the traffic outside the tunnel. This means that although either IPv4 or IPv6 may be carried inside of the tunnel, to use IPv6 traffic inside the tunnel it must be connected between IPv6 peer IP addresses, not IPv4. In other words, the inner and outer address family must match, they cannot be mixed.

As with most other shortcomings of IKEv1, this has been addressed in IKEv2. Tunnels using IKEv2 may carry both types of traffic no matter which protocol is used to establish the outer tunnel. With IKEv2, mobile clients may also use both IPv4 and IPv6, provided the client supports it.

See also:

- *IPsec Logs*
- *IPsec Status*
- *IPsec Site-to-Site VPN Example with Pre-Shared Keys*
- *IPsec Site-to-Site VPN Example with Certificate Authentication*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-RADIUS*
- *IPsec Remote Access VPN Example Using IKEv2 with EAP-TLS*
- *IPsec Remote Access VPN Example Using IKEv1 with Xauth*
- *IPsec Remote Access VPN Example Using IKEv1 with Pre-Shared Keys*
- *Routing Internet Traffic Through a Site-to-Site IPsec Tunnel*
- *Connecting to Cisco IOS Devices with IPsec*
- *Connecting to Cisco PIX/ASA Devices with IPsec*
- *Connecting to L2TP/IPsec from Android*
- *L2TP/IPsec Remote Access VPN Configuration Example*

- *Troubleshooting IPsec VPNs*

IPsec provides a standards-based VPN implementation that is compatible with a wide range of clients for mobile connectivity, and other firewalls and routers for site-to-site connectivity. It supports numerous third party devices and is being used in production with devices ranging from consumer grade Linksys routers all the way up to IBM z/OS mainframes, and everything imaginable in between. This chapter describes the configuration options available, and how to configure various common scenarios.

See also:

For general discussion of the various types of VPNs available in AZTCO-FW® software and their pros and cons, see *Virtual Private Networks*.

AZTCO-FW software supports IPsec with IKEv1 and IKEv2, multiple phase 2 definitions for each tunnel, as well as NAT traversal, NAT on Phase 2 definitions, a large number of encryption and hash options, and many more options for mobile clients, including xauth and EAP.

14.7.14 IPsec Terminology

Before delving too deeply into configuration, there are a few terms that are used throughout the chapter that require explanation. Other terms are explained in more detail upon their use in configuration options.

IKE

IKE stands for *Internet Key Exchange*, and comes in two different varieties: IKEv1 and IKEv2. Nearly all devices that support IPsec use IKEv1. A growing number of devices also support the newer IKEv2 protocol which is an updated version of IKE that solves some of the difficulties present in the earlier version. For example, IKEv2 has MOBIKE, which is a standard for mobile clients that allows them to switch addresses dynamically. It also has built-in NAT traversal, and standard mechanisms for reliability similar to DPD. In general IKEv2 provides a more stable and reliable experience, provided both ends support it sufficiently.

ISAKMP Security Association

ISAKMP stands for Internet Security Association and Key Management Protocol. It gives both parties a mechanism by which they can set up a secure communications channel, including exchanging keys and providing authentication.

An ISAKMP Security Association (ISAKMP SA) is a one-way policy which defines how traffic will be encrypted and handled. Each active IPsec tunnel will have two security associations, one for each direction. The ISAKMP Security Associations are setup between the public IP addresses for each endpoint. Knowledge of these active security associations is kept in the Security Association Database (SAD).

Security Policy

A Security Policy manages the complete specifications of the IPsec tunnel. As with Security Associations, these are one-way, so for each tunnel there will be one in each direction. These entries are kept in the Security Policy Database (SPD). The SPD is populated with two entries for each tunnel connection as soon as a tunnel is added. By contrast, SAD entries only exist upon successful negotiation of the connection.

In AZTCO-FW software, Security Policies control which traffic will be intercepted by the kernel for delivery via IPsec.

Phase 1

There are two phases of negotiation for an IPsec tunnel. During phase 1, the two endpoints of a tunnel setup a secure channel between using ISAKMP to negotiate the SA entries and exchange keys. This also includes authentication, checking identifiers, and checking the pre-shared keys (PSK) or certificates. When phase 1 is complete the two ends can exchange information securely, but have not yet decided what traffic will traverse the tunnel or its encryption.

Phase 2

In phase 2, the two endpoints negotiate how to encrypt and send the data for the private hosts based on Security Policies. This part builds the tunnel used for transferring data between the endpoints and clients whose traffic is handled by those endpoints. If the policies on both side agree and phase 2 is successfully established, the tunnel will be up and ready for use for traffic matching the phase 2 definitions.

See also:

- *L2TP Server Configuration*
- *Troubleshooting Cisco VPN Pass Through*

VPNs provide a means of tunneling traffic through an encrypted connection, preventing it from being seen or modified in transit. AZTCO-FW® software offers three VPN options: IPsec, OpenVPN, and L2TP. This chapter provides an overview of VPN usage, the pros and cons of each type of VPN in AZTCO-FW, and how to decide which is the best fit for a particular environment. Subsequent chapters discuss each VPN option in detail.

L2TP is purely a tunneling protocol and does not offer any encryption of its own. It is typically combined with some other method of encryption such as IPsec in transport mode. Because of this, it doesn't fit in with most of the discussion in this chapter. See *L2TP VPN* for more information on L2TP.

14.8 PPTP Warning

PPTP server support has been removed from AZTCO-FW software. Despite the attraction of its convenience, PPTP must not be used under any circumstances because it is no longer secure. This is not specific to the implementation of PPTP that was in AZTCO-FW; Any system that handles PPTP is no longer secure. The reason for the insecurity is that PPTP relies upon MS-CHAPv2 which has been completely compromised. Intercepted traffic can be decrypted by a third party 100% of the time, so consider any traffic carried in PPTP unencrypted. Migrate to another VPN type such as OpenVPN or IPsec as soon as possible. More information on the PPTP security compromise can be found at https://isc.sans.edu/diary/End+of+Days+for+MS-CHAPv2/13807 and https://www.cloudcracker.com/blog/2012/07/29/cracking-ms-chap-v2/.

14.9 Common deployments

There are four common uses of the VPN capabilities of AZTCO-FW, each covered in this section.

14.9.1 Site-to-site connectivity

Site-to-site connectivity is primarily used to connect networks in multiple physical locations where a dedicated, always-on, connection between the locations is required. This is frequently used to connect branch offices to a main office, connect the networks of business partners, or connect a network to another location such as a data center environment.

Before the proliferation of VPN technology, private WAN circuits were the only solution to connect multiple locations.
These technologies include point-to- point dedicated circuits, packet switching technologies such as frame relay and ATM, and more recently, MPLS (Multiprotocol Label Switching) and fiber and copper based metropolitan Ethernet services. While these types of private WAN connectivity provide reliable, low latency connections, they are also very costly with recurring monthly fees. VPN technology has grown in popularity because it provides the same secure site to site connectivity using Internet connections that are generally much less costly.

Limitations of VPN connectivity

Performance is an important consideration when planning a VPN solution. In some networks, only a private WAN circuit can meet the requirements for bandwidth or latency. Latency is usually the biggest factor. A point to point DS1 circuit has end to end latency of about 3-5 ms, while the latency to the first hop on an ISP network will generally be at least that much if not higher. Metro Ethernet services or fiber circuits have end to end latency of about 0-3 ms, usually less than the latency to the first hop of an ISP network. That will vary some based on geographical distance between the sites. The stated numbers are typical for sites within a couple hundred miles of each other. VPNs usually see latency of around 30-60 ms depending on the Internet connections in use and the geographical distance between the locations. Latency can be minimized and VPN performance maximized by using the same ISP for all VPN locations, but this isn’t always feasible.

Certain protocols perform very poorly with the latency inherent in connections over the Internet. Microsoft file sharing (SMB) is a common example. At sub-10 ms latency, it performs well. At 30 ms or higher, it’s sluggish, and at more than 50 ms it’s painfully slow, causing frequent hangs when browsing folders, saving files, etc. Getting a simple directory listing requires numerous round trip connections between the client and server, which significantly exacerbates the increased delay of the connection. In Windows Vista and Server 2008, Microsoft introduced SMB 2.0 which includes new capabilities to address the issue described here. SMB 2.0 enables the sending of multiple actions in a single request, as well as the ability to pipeline requests, meaning the client can send additional requests without waiting for the response from prior requests. If a network uses exclusively Vista and Server 2008 or newer operating systems this won’t be a concern, but given the rarity of such environments, this will usually be a consideration. SMB 3.0 further improves in this area with support for multiple streams.

Two more examples of latency sensitive protocols are Microsoft Remote Desktop Protocol (RDP) and Citrix ICA. There is a clear performance and responsiveness difference with these protocols between sub-20 ms response times typically found in a private WAN, and the 50-60+ ms response times common to VPN connections. If remote users work on published desktops using thin client devices, there will be a notable performance difference between a private WAN and VPN. Whether that performance difference is significant enough to justify the expense of a private WAN will vary from one environment to another.

There may be other network applications in an environment that are latency sensitive, where the reduced performance of a VPN is unacceptable. Or all locations may be within a relatively small geographical area using the same ISP, where the performance of a VPN rivals that of private WAN connections.

14.9.2 Remote access

Remote access VPNs enable users to securely connect into a network from any location where an Internet connection is available. This is most frequently used for mobile workers (often referred to as "Road Warriors") whose job requires frequent travel and little time in the office, and to give employees the ability to work from home. It can also allow contractors or vendors temporary access to a network. With the proliferation of smart phones, users have a need to securely access internal services from their phones using a remote access VPN.

14.9.3 Protection for wireless networks

A VPN can provide an additional layer of protection for wireless networks. This protection is two-fold: It provides an additional layer of encryption for traffic traversing the wireless network, and it can be deployed in such a way that it requires additional authentication before access to network resources is permitted. This is deployed mostly the same as remote access VPNs. This is covered in *Additional protection for a wireless network*.

14.9.4 Secure relay

Remote access VPNs can be configured in a way that passes all traffic from the client system over the VPN. This is nice to have when using untrusted networks, such as wireless hotspots as it lets a client push all its Internet traffic over the VPN and out to the Internet from the VPN server. This protects the user from a number of attacks that are possible on untrusted networks, though it does have a performance impact since it adds additional hops and latency to all connections. That impact is usually minimal with high speed connectivity when the client and VPN server are relatively close geographically.

CHAPTER

FIFTEEN

L2TP VPN

15.1 L2TP and Firewall Rules

By default, when the L2TP server is enabled, firewall rules will not be automatically added to the chosen interface to permit UDP port *1701*. A firewall rule must be added to whichever interface the L2TP traffic will be entering, typically WAN, the WAN containing the default gateway, or IPsec.

15.2 L2TP and Multi-WAN

L2TP uses UDP port 1701. Because L2TP relies on UDP, the server may have issues using any WAN that is not the default gateway. The daemon will respond from the firewall using the closest address to the client, following the routing table, which is the WAN with the default gateway for remote clients.

15.3 L2TP Server Configuration

To use L2TP, first browse to VPN > L2TP. Select Enable L2TP server.

Warning: L2TP is not a secure protocol by itself; it only provides tunneling, it does not perform encryption.

See also:

L2TP/IPsec is a way to secure L2TP traffic by sending it through an encrypted IPsec tunnel. This may be used in combination with a mobile IPsec setup to configure L2TP+IPsec; see *L2TP/IPsec Remote Access VPN Configuration Example* for details.

15.3.1 Interface

The Interface setting controls where the L2TP daemon will bind and listen for connections. This is typically the *WAN* interface accepting inbound connections.

15.3.2 IP Addressing

Before starting, determine which IP addresses to use for the L2TP server and clients and now many concurrent clients to support.

Server Address An *unused* IP address outside of the Remote Address Range, such as 10.3.177.1 as shown in Figure *L2TP IP Addressing*.

Remote Address Range Usually a new and unused subnet, such as 10.3.177.128/25 (.128 through .255). These are the addresses to be assigned to clients when they connect.

Number of L2TP users Controls how many L2TP users will be allowed to connect at the same time, in this example *13* has been selected.

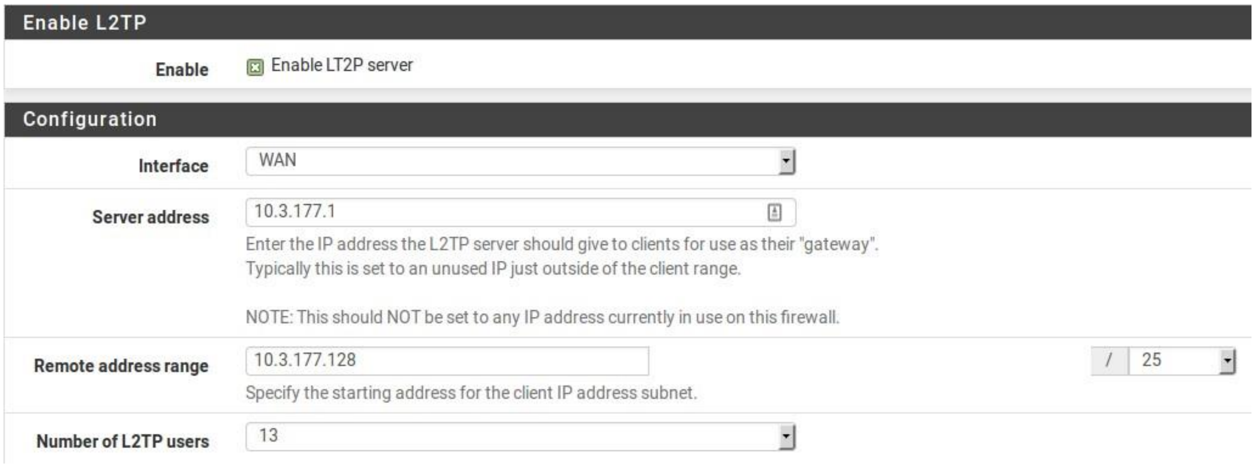

Enable L2TP

Enable ☒ Enable LT2P server

Configuration

Interface: WAN

Server address: 10.3.177.1

Enter the IP address the L2TP server should give to clients for use as their "gateway".
Typically this is set to an unused IP just outside of the client range.

NOTE: This should NOT be set to any IP address currently in use on this firewall.

Remote address range: 10.3.177.128 / 25

Specify the starting address for the client IP address subnet.

Number of L2TP users: 13

Fig. 1: L2TP IP Addressing

DNS servers can also be defined for end users when needed. Fill in the Primary and Secondary L2TP DNS server fields with the DNS server IP addresses for connecting clients.

15.3.3 Authentication

Secret Required by some L2TP implementations, similar to a group password or pre-shared key. Support for this varies from client to client. Leave the field blank unless it is known to be required. If required, enter and confirm the secret.

Authentication Type Decides between *PAP*, *CHAP*, or *MS-CHAPv2* authentication for users. Support for this can vary from client to client and it may also depend on the RADIUS server as well. The *CHAP* based types are more secure, but *PAP* is more widely compatible.

Users may be authenticated from the local user database, or via an external RADIUS server. This can be used to authenticate L2TP users from Microsoft Active Directory (see *Authenticating from Active Directory using RADIUS/NPS*) as well as numerous other RADIUS capable servers.

If using RADIUS, check the Use a RADIUS server for authentication box and fill in the RADIUS server and shared secret. For authentication using the local user database, leave that box unchecked. Users must be added manually on the Users tab of the VPN > L2TP screen unless using RADIUS. See *Adding Users* below for more details on the built-in authentication system.

20.3.4 Save changes to start L2TP server

After filling in the aforementioned items, click Save. This will save the configuration and launch the L2TP server.

15.3.5 Configure firewall rules for L2TP clients

Browse to Firewall > Rules and click the L2TP VPN tab. These rules control traffic from L2TP clients. Until a firewall rule has been added to allow traffic, all traffic initiated from connected L2TP clients will be blocked. Traffic initiated from the LAN to L2TP clients is controlled using LAN firewall rules. Initially an allow all rule may be desired here for testing purposes as shown in Figure *L2TP VPN Firewall Rule*, and once functionality has been verified, restrict the ruleset as desired.

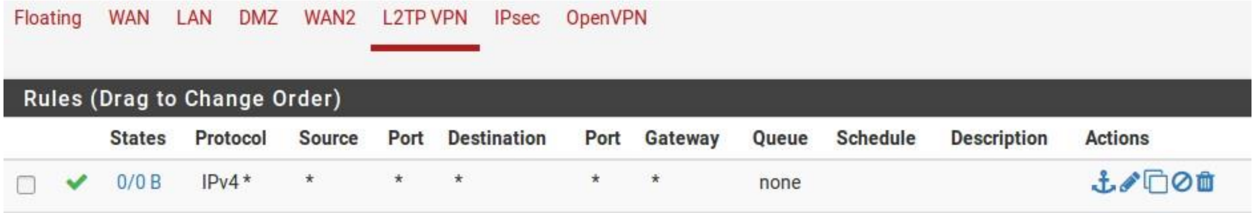

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
□	✓	0/0 B	IPv4 *	*	*	*	*	*	none			

Fig. 2: L2TP VPN Firewall Rule

Note: Remember that a rule must also be added to the interface receiving the L2TP traffic, typically WAN or IPsec, to pass UDP to the firewall with a destination port of 1701.

15.3.6 Adding Users

Adding users to the built-in L2TP users system is simple. To add local users:

- Navigate to VPN > L2TP, Users tab. The users screen as shown in Figure *L2TP Users Tab* will be presented.
- Click ➕ Add to show the form used to add users.

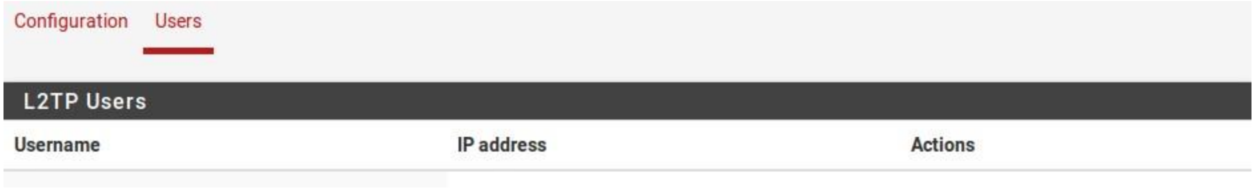

L2TP Users

Username	IP address	Actions

Fig. 3: L2TP Users Tab

- Enter the Username, Password and Confirm Password for a user, as in Figure *Adding a L2TP User*.
- Enter a static IP assignment if desired.
- Click Save, and then the user list will return.
- Repeat the process for each user to add.

To edit an existing user, click ⧉. Users may be deleted by clicking 🗑.

20.3. L2TP Server Configuration

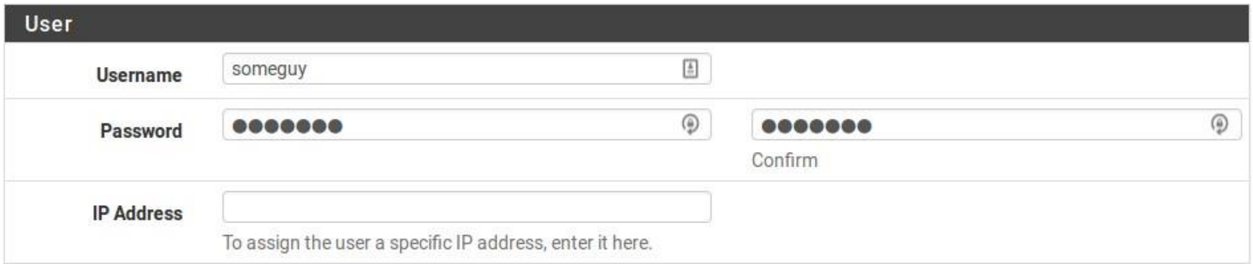

User

Username	someguy	
Password	●●●●●●●	●●●●●●● Confirm
IP Address	To assign the user a specific IP address, enter it here.	

Fig. 4: Adding a L2TP User

See also:

- *L2TP/IPsec Remote Access VPN Configuration Example*
- *Troubleshooting L2TP*
- *L2TP Logs*

AZTCO-FW® software can act as an L2TP VPN server. L2TP is purely a tunneling protocol that offers no encryption of its own, so it is typically combined with some other encryption technique, such as IPsec.

Warning: AZTCO-FW supports L2TP/IPsec, however, some clients will not work properly in many common scenarios. The most common problem scenario is Windows clients behind NAT, in that case the Windows client and the strongSwan IPsec daemon are not fully compatible, which leads to failure. In these situations, we recommend using IKEv2 instead.

See also:

IPsec Remote Access VPN Example Using IKEv2 with EAP-MSCHAPv2 contains a walkthrough for configuring IKEv2, which is a much more flexible solution.

See also:

For general discussion of the various types of VPN implementations available in AZTCO-FW and their pros and cons, see *Virtual Private Networks*.

15.4 L2TP Security Warning

L2TP on its own is not encrypted, so it is not intended for private traffic. Some devices, such as Android, offer an L2TP-only client which is capable of connecting back to AZTCO-FW but it should only be used for traffic that is already encrypted, or if the traffic is not considered private. For example, tunneling Internet traffic so it appears to originate from another location.

CHAPTER

SIXTEEN

SERVICES

16.1 DHCPv4 Server

To alter the behavior of the IPv4 DHCP server, navigate to Services > DHCP Server in the web interface. The behavior of the IPv4 DHCP server is controlled there, along with static IP address mappings and related options such as static ARP.

16.1.1 Choosing an Interface

The DHCP configuration page contains a tab for each interface with a static IP address. Each interface has its own separate DHCP server configuration, and they may be enabled or disabled independently of one another. Before making any changes, visit the tab for the correct interface.

16.1.2 General Options

Enable The first setting on the tab enables or disables DHCP service for the interface. To turn on DHCP for the interface, check Enable DHCP server on [name] interface. To disable the service, uncheck the box instead.

Deny unknown clients Under normal circumstances, the DHCP server will answer requests from any client requesting a lease. In most environments this is normal and acceptable behavior, but in restricted or secure environments this behavior is undesirable. With this option set, only clients with static mappings defined will receive leases. This is a more secure practice but is much less convenient. This option is per-pool, meaning that if unknown clients are denied in the default range, another pool of IP addresses may be defined that does not have the setting checked. The DHCP server will assign clients IP addresses out of that alternate pool instead.

> Note: This will protect against low-knowledge users and people who casually plug in devices. Be aware, however, that a user with knowledge of the network could hardcode an IP address, subnet mask, gateway, and *DNS* which will still give them access. They could also alter/spoof their MAC address to match a valid client and still obtain a lease. Where possible, couple this setting with static ARP entries, access control in a switch that will limit MAC addresses to certain switch ports for increased security, and turn off or disable unused switch ports.

Subnet The network address of the interface subnet, for reference purposes.

Subnet Mask The subnet mask for the interface subnet, for reference purposes.

Available Range The range of available addresses inside the interface subnet, for reference and to help determine the desired range for DHCP clients. The network address and broadcast address are excluded, but interface addresses and Virtual IP addresses are not excluded.

Range This defines the DHCP address range, also referred to as the Scope or Pool. The two boxes for Range tell the firewall the first and last address for use as a DHCP pool. Addresses between the entered values, inclusive, will be used for clients which request addresses via DHCP. The range must be entered with the lower number first, followed by the higher number. For example, the default LAN DHCP range is based off of the subnet for the default LAN IP address. It is 192. 168.1.100 to 192.168.1.199. This range can be as large or as small as the network needs, but it must be wholly contained within the subnet for the interface being configured.

16.1.3 Additional Pools

The Additional Pools section defines extra pools of addresses inside of the same subnet. These pools can be used to craft sets of IP addresses specifically for certain clients, or for overflow from a smaller original pool, or to split up the main pool into smaller chunks with a GAP of non-DHCP IP addresses in the middle of what used to be the pool. A combination of the MAC Address Control options may be used to guide clients from the same manufacturer into a specific pool, such as VoIP phones.

To add a new pool, click + Add Pool and the screen will switch to the pool editing view, which is nearly the same as the normal DHCP options, except a few options that are not currently possible in pools are omitted. The options behave the same as the others discussed in this section. Items left blank will, by default, fall through and use the options from the main DHCP range.

Note: See the MAC Address Control section below for specifics on directing clients into or away from pools.

16.1.4 Servers

WINS Servers Two WINS Servers (Windows Internet Name Service) may be defined that will be passed on to clients. If one or more WINS servers is required, enter their IP addresses here. The actual servers do not have to be on this subnet, but be sure that the proper routing and firewall rules are in place to let them be reached by client PCs. If this is left blank, no WINS servers will be sent to the client.

DNS Servers The DNS Servers may or may not need filled in, depending on the firewall configuration. If the built-in DNS Resolver or DNS Forwarder is used to handle DNS, leave these fields blank and AZTCO-FW® will automatically assign itself as the DNS server for client PCs. If the DNS forwarder is disabled and these fields are left blank, AZTCO-FW will pass on whichever DNS servers are defined under System > General Setup. To use custom DNS Servers instead of the automatic choices, fill in the IP addresses for up to four DNS servers here. In networks with Windows servers, especially those employing Active Directory, it is recommended to use those servers for client DNS. When using the DNS Resolver or DNS forwarder in combination with CARP, specify the CARP Virtual IP address on this interface here.

16.1.5 Other Options

Gateway This may also be left blank if this firewall is acting as the gateway for the network on this interface. If that is not the case, fill in the IP address for the gateway to be used by clients on this interface. When using CARP, fill in the CARP Virtual IP address on this interface here.

Domain Name Specifies the domain name passed to the client to form its fully qualified hostname. If the Domain Name is left blank, then the domain name of the firewall it sent to the client. Otherwise, the client is sent this value.

Domain Search List Controls the DNS search domains that are provided to the client via DHCP. If multiple domains are present and short hostnames are desired, provide a list of domain names here, separated by a semicolon. Clients will attempt to resolve hostnames by adding the domains, in turn, from this list before trying to find them externally. If left blank, the Domain Name option is used.

Note: The Domain Search List is provided via DHCP option 119. As of this writing, no Windows DHCP *client* of any version supports DHCP option 119. Other operating systems such as BSD, Linux, and OS X do support obtaining the Domain Search List via DHCP option 119.

Default lease time Controls how long a lease will last when a client does not request a specific lease length. Specified in seconds, default value is 7200 seconds (2 hours)

Maximum lease time Limits a requested lease length to a stated maximum amount of time. Specified in seconds, default value is 86400 seconds (1 day).

Failover Peer IP If this system is part of a High Availability failover cluster, enter the real IP address of the other system in this subnet here. Do not enter a CARP Virtual IP address.

Static ARP This checkbox works similar to denying unknown MAC addresses from obtaining leases, but takes it a step further in that it also restricts any unknown MAC address from communicating with this firewall. This stops would- be abusers from hardcoding an unused address on this subnet, circumventing DHCP restrictions.

Note: When using static ARP, all systems that need to communicate with the firewall must be listed in static mappings before activating this option, especially the system being used to connect to the AZTCO-FW GUI. Also be aware that this option may prevent people from hardcoding an IP address and talking to the firewall, but it does not prevent them from reaching each other on the local network segment.

Time Format Change By default, the ISC DHCP daemon maintains lease times in UTC. When this option is checked, the times on the DHCP Leases status page are converted to the local time zone defined on the firewall.

Statistics Graphs This option, disabled by default, activates RRD graphing for monitoring the DHCP pool utilization.

16.1.6 Dynamic DNS

For Dynamic DNS settings, click Display Advanced to the right of that field, which displays the following options:

Enable Check the box to enable registration of DHCP client names in DNS using an external (non-AZTCO-FW) DNS server.

DDNS Domain The domain name used for registering clients in DNS

Primary DDNS Address The DNS server used for registering clients in DNS

DNS Domain Key The encryption key used for DNS registration

DNS Domain Key Secret The secret for the key used for DNS registration

16.1.7 MAC Address Control

For MAC Address Control, click Display Advanced to show the lists of allowed and denied client MAC addresses. Each list is comma-separated and contains portions of MAC addresses. For example, a group of VoIP phones from the same manufacturer may all start with the MAC address aa:bb:cc. This can be leveraged to give groups of devices or users separate DHCP options.

Allow A list of MAC Addresses to allow in this pool. If a MAC address is in the allow box, then all others will be denied except the MAC address specified in the allow box.

Deny A list of MAC Addresses to deny from this pool. If a MAC address is in the deny list, then all others are allowed.

It is best to use a combination of allow and deny to get the desired result, such as: In the main pool, leave allow blank and deny aa:bb:cc. Then in the VoIP pool, allow aa:bb:cc. If that extra step is not taken to allow the MAC prefix in the additional pool, then other non-VoIP phone clients could receive IP addresses from that pool, which may lead to undesired behavior.

This behavior may also be used to blacklist certain devices from receiving a DHCP response. For example to prevent Example brand printers from receiving a DHCP address, if MAC addresses all start with ee:ee:ee, then place that in the deny list of each pool.

16.1.8 NTP Servers

To specify NTP Servers (Network Time Protocol Servers), click the Display Advanced button to the right of that field, and enter IP addresses for up to two NTP servers.

16.1.9 TFTP Server

click the Display Advanced button next to TFTP to display the TFTP server option. The value in the TFTP Server box, if desired, must be an IP address or hostname of a TFTP server. This is most often used for VoIP phones, and may also be referred to as “option 66” in other documentation for VoIP and DHCP.

16.1.10 LDAP URI

click the Display Advanced button next to LDAP to display the LDAP Server URI option. LDAP Server URI will send an LDAP server URI to the client if requested. This may also be referred to as DHCP option 95. It takes the form of a fully qualified LDAP URI, such as ldap://ldap.example.com/dc=example,dc=com. This option can help clients using certain kinds of systems, such as OpenDirectory, to find their server.

16.1.11 Additional BOOTP/DHCP Options

Other numeric DHCP options can be sent to clients using the Additional BOOTP/DHCP Options controls. To view these options, click Display Advanced in this section. To add a new option, click + Add.

Number The DHCP option code number. IANA maintains a list of all valid DHCP options.

Type The choices and formats for each type may be a little counter-intuitive, but the labels are used directly from the DHCP daemon. The proper uses and formats are:

Text Free-form text to be sent in reply, such as http://www.example.com/wpad/ wpad.dat or Example Company.

String A string of hexadecimal digits separated by a colon, such as c0:a8:05:0c.

Boolean Either true or false.

Unsigned 8, 16, or 32-bit Integer A positive Integer that will fit within the given data size, such as 86400.

Signed 8, 16, or 32-bit Integer A positive or negative Integer that will fit within the given data size, such as -512.

IP address or host An IP address such as 192.168.1.1 or a hostname such as www. example.com.

Value The value associated with this numeric option and type.

For more information on which options take a specific type or format, see the linked list above from the IANA.

Note: When using numbered custom options, be careful of the type. Some will be OK on text/string but others are not.

For example, DHCP options for code 132 (and presumably 133) for VLAN ID must be set for a type of unsigned integer 32.

16.1.12 Network Booting

To view the Network boot settings, click ⊕ in the Network Booting section header bar.

Enable Check to enable network booting options in DHCP

Next Server The IP address from which boot images are available

Default BIOS file name File name for the boot image (Non-UEFI)

UEFI 32 bit file name File for 32-bit UEFI booting

UEFI 64 bit file name File for 64-bit UEFI booting

Root Path String to target a specific device as the client's root filesystem device, such as iscsi:(servername):(protocol):(port):(LUN):targetname.

16.1.13 Save Settings

After making changes, click Save before attempting to create static mappings. Changes to settings will be lost if the browser leaves this page without saving.

16.1.14 Static Mappings

Static DHCP mappings express a preference for which IP address will be assigned to a given client based on its MAC address. In a network where unknown clients are denied, this also serves as a list of "known" clients which are allowed to receive leases or have static ARP entries. Static mappings can be added in one of two ways:

- From this screen, click + Add.
- Add them from the DHCP leases view, which is covered later in this chapter.

On this screen, only the MAC address is necessary.

MAC Address The client MAC address which identifies the host to deliver options on this page, or by entering only the MAC address, it will be added to the list of known clients for use when the Deny unknown clients option is set.

> Note: Client MAC address can be obtained from a command prompt on most platforms. On UNIX-based or UNIX-work-alike operating systems including Mac OS X, typing ifconfig -a will show the MAC address for each interface. On Windows-based platforms, ipconfig /all will show the MAC address. The MAC address may also sometimes be found upon a sticker on the network card, or near the network jack for integrated adapters. For hosts on the same subnet, the MAC can be determined by pinging the IP address of the host and then running arp - a.

Client Identifier An ID sent by the client to identify itself.

IP Address The IP address field is needed if this will be a static IP address mapping instead of only informing the DHCP server that the client is valid. This IP address is a preference, not a reservation. Assigning an IP address here will not prevent someone else from using the same IP address. If this IP address is in use when this client requests a lease, it will instead receive an address from the general pool. For this reason, the AZTCO-FW WebGUI does not allow assigning static IP mappings inside of the DHCP pool.

Hostname The hostname of the client. This does not have to match the actual hostname set on the client. The hostname set here will be used when registering DHCP addresses in the DNS forwarder.

Description Cosmetic only, and available for use to help track any additional information about this entry. It could be the name of the person who uses the PC, its function, the reason it needed a static address, or the administrator who added the entry. It may also be left blank.

ARP Table Static Entry If checked, this entry will receive a static ARP entry in the OS tying this IP address to this MAC address.

Note: If this option is used rather than using the global static ARP option, it does not prevent that MAC address from using other IP addresses, it only prevents other MAC addresses from using this

IP address. In other words, it prevents another machine from using that IP to reach the firewall, but it doesn't stop the user from changing their own IP address to something different.

The remaining options available to set for this client are the same in behavior to the ones found earlier in this section for the main DHCP settings.

Click Save to finish editing the static mapping and return to the DHCP Server configuration page.

16.2 DHCPv6 Server

The DHCPv6 server in AZTCO-FW® software will hand out addresses to DHCPv6 clients and automatically configure them for network access. By default, the DHCPv6 server is enabled on the LAN interface and set to use a prefix obtained by tracking WAN's DHCPv6 delegation.

The DHCPv6 server page, found under Services > DHCPv6 Server, has a tab for each available interface. The DHCPv6 daemon can only run and be configured on interfaces with a Static IP address, so if a tab for an interface is not present, check that it is enabled and set with a Static IP. It is not currently possible to adjust settings for tracked interface DHCP service.

The DHCPv6 server cannot be active on any interface if the *DHCPv6 Relay* service is in use.

16.2.1 DHCP Instance Options

For each Interface, there are many options to choose from. At a minimum, the Enable box must be checked on the interface tab and an address range (starting and ending IPv6 addresses) to use for DHCPv6 clients must be defined. For the DHCPv6 server to be active on the network, *Router Advertisements* must also be set to either *Managed* or *Assisted* mode on the Router Advertisements tab.

The other settings may be configured, but are optional.

Note: DHCPv6 does not provide gateway information. *Router Advertisements* tell hosts on the network how to reach a router. DHCPv6 is for other host configuration such as DNS, delegation, and so on.

See the *DNS Forwarder* article for information on the default DNS server behavior.

Some other options which may be set for clients include Network booting options, LDAP URI, and the ability to add in any custom DHCP option number and value.

16.2.2 DHCPv6 Range

The Range parameter works similarly to the same setting on IPv4 but it is worth mentioning again here due to the differences in IPv6 addressing.

Given the vast amount of space available inside even a /64, a good trick is to craft a range that restricts hosts to use an easy to remember or recognize range. For example, Inside a /64 such as 2001:db8:1:1::, set the DHCPv6 range be: 2001:db8:1:1::d:0000 to 2001:db8:1:1::d:FFFF, using the d in the second to last section of the address as a sort of shorthand for "DHCP". That example range contains 2^16 (65,536) IPs, which is extremely large by today's IPv4 standards, but only a small portion of the whole /64.

16.2.3 DHCPv6 Prefix Delegation

Prefix delegation, covered earlier in *DHCP6 Prefix Delegation* and *Track Interface*, allows automatically dividing and allocating a block of IPv6 addresses to networks that will live behind other routers and firewall that reside downstream from AZTCO-FW (e.g. in the LAN, DMZ, etc). Most users acting in a client capacity will not need this and will likely leave it blank.

Prefix delegation can be used to hand out /64 chunks of a /48 to routers automatically, or any other combination, so long as the range is set on the boundaries of the desired delegation size. The downstream router obtains an IPv6 address and requests a delegation, and the server allocates one and dynamically adds a route so that it is reachable via the assigned DHCPv6 address given to the client.

The Prefix Delegation Range Sets the start and end of the delegation pool. The range of IPv6 addresses specified here must be routed to this firewall by upstream routers. For example, to allocate /60 networks to downstream firewalls out of a given range, then one could specify 2001:db8:1111:F000:: to 2001:db8:1111:FFF0:: with a Prefix Delegation Size of *60*. This allocates a /60 (16 subnets of size /64) to each downstream firewall that requests a delegation so that they can in turn use those for their LAN, VPNs, DMZ, etc. Downstream firewalls can even further delegate their own allocation to routers behind them. Note that in this example, 16 delegations would be possible. Adjust the range and size as needed.

When crafting the values for the range and delegation size, keep in mind that the range must start and end on boundaries that align with the desired prefix size. In this /60 example, the range could not start or end on anything that has a value in the places to the right of the second value in the fourth section of the address, so it can start on 2001:db8:1111:F500:: but not 2001:db8:1111:F550::.

16.2.4 DHCPv6 Static Mappings

Static mappings on DHCPv6 work differently than IPv4. On IPv4, the mappings were performed using the MAC address of the PC. For IPv6, the designers decided that wasn't good enough, since the MAC address of a PC could change, but still be the same PC.

Enter, the DHCP Unique Identifier, or DUID. The DUID of the host is generated by the operating system of the client and, in theory, will remain unique to that specific host until such time as the user forces a new DUID or the operating system is reinstalled. The DUID can range from 12 to 20 bytes, and varies depending on its type.

The DUID field on the static mapping page expects a DUID for a client PC in a special format, represented by pairs of hexadecimal digits, separated by colons, such as 00:01:00:01:1b:a6:e7:ab:00:26:18:1a:86:21.

How to obtain this DUID depends on the operating system. The easiest way is to allow the PC to obtain a lease via DHCPv6, and then add an entry from the DHCPv6 Leases View (Status DHCPv6 Leases). In Windows, it can be found as DHCPv6 Client DUID in the output of ipconfig /all.

Note: On Windows, the DUID is generated at install time, so if a base image is used and workstations are cloned from there, they can all end up with the same DUID, and thus all end up pulling the same IPv6 address over DHCPv6.

Clear the DUID from the registry before making an image to clone, by issuing the following command:

```
reg delete HKLM\SYSTEM\CurrentControlSet\Services\Tcpip6\Parameters /f /v Dhcpv6DUID
```

That command may also be run on a working system to reset its DUID if needed.

16.2. DHCPv6 Server

DUID Format

The DUID format is listed on the page, but it roughly follows the format:

```
DUID-LLT - ETH -- TIME --- ---- address ----
```

DUID-LLT is link-layer plus time, which means it uses the link type of a network interface on the system (Generally 00:01 to indicate the format, plus 00:01, or 00:06 for Ethernet), plus the timestamp at which the DUID was generated in hex, plus the MAC address of the first NIC. It may be difficult or impossible to predict a system’s DUID. Unless the operating system has a way to look it up, it may be best to allow the client to obtain a dynamic lease and then copy the DUID from the leases view.

16.2.5 Numbered Options Notes

When using numbered custom options, be careful of the type. Some will be OK on text/string but others are not. Also beware that numbered options do NOT correspond exactly to the DHCP numbered options for IPv4

For more information on DHCP option numbers and types, see https://tools.ietf.org/html/draft-ietf-dhc-v6opts-00

16.3 IPv6 Router Advertisements

Automatic address assignment for IPv6 works quite a bit differently than IPv4. Even so, most of the DHCP options are similar, but there are notable differences in behavior in how things are assigned and also how items like the gateway are handed off to clients. Unless otherwise noted, options of the same name work the same for DHCP and

DHCPv6. DHCPv6 and Router Advertisements (RA) are configured under Services > DHCPv6 Server/RA. Under that page there are two tabs: One for DHCPv6 Server and one for Router Advertisements.

16.3.1 DHCPv6 vs Stateless Address Autoconfiguration

There are a few clients that do not have support for DHCPv6. Some clients only support Stateless Address Autoconfiguration, or SLAAC for short. There is no way for the firewall to have direct knowledge of a list of hosts on the segment using SLAAC addresses, so for some environments it is much less desirable because of the lack of control and reporting of addresses. Consider address tracking and operating system support requirements when deciding how to allocate IPv6 addresses to clients on the network.

Many operating systems such as Windows, OS X, FreeBSD, Linux, and their cousins contain DHCPv6 clients that are capable of obtaining addresses as expected via DHCPv6. Some lightweight or mobile operating systems such as Android do not contain a DHCPv6 client and will only function on a local segment with IPv6 using SLAAC.

16.3.2 Router Advertisements (Or: "Where is the DHCPv6 gateway option")

In IPv6, a router is located through Router Advertisement (RA) messages sent from routers instead of by DHCP; IPv6-enabled routers that support dynamic address assignment are expected to announce themselves on the network to all clients. As such, DHCPv6 does not include any gateway information. So clients can obtain their addresses from DHCPv6 or SLAAC, but unless they are statically configured, they always locate their next hop by using RA packets sent from available gateways.

To enable the RA service:

- Navigate to Services > DHCPv6 Server/RA
- Click the interface tab for the interface being configured

21.3. IPv6 Router Advertisements

- Click the Router Advertisements tab
- Select a mode other than *Disabled* from the Router Mode drop-down list
- Click Save

The other options to control RA behavior may be set as needed for the network:

Router Advertisement Modes The modes for the RA daemon control the services offered by AZTCO-FW®, announce the firewall as an IPv6 router on the network, and direct clients on how to obtain addresses.

Disabled The RA daemon is disabled and will not run. IPv6 gateways must be entered manually on any client hosts.

Router Only This firewall will send out RA packets that advertise itself as an IPv6 router. DHCPv6 is disabled in this mode.

Unmanaged The firewall will send out RA packets and clients are directed to assign themselves IP addresses within the interface subnet using SLAAC. DHCPv6 is disabled in this mode.

Managed The firewall will send out RA packets and addresses will only be assigned to clients using DHCPv6.

Assisted The firewall will send out RA packets and addresses can be assigned to clients by DHCPv6 or SLAAC.

Stateless DHCP The firewall will send out RA packets and addresses can be assigned to clients by SLAAC while providing additional information such as DNS and NTP from DHCPv6.

Router Priority If multiple IPv6 routers exist on the same network segment, they can indicate to clients in which order they should be used. If a high priority router becomes unavailable, clients will try a normal priority router, and finally a low priority router. Select either *Low*, *Normal*, or *High* from the list. If there is only one router on the network, use *Normal*.

Default Valid Lifetime Length of time, specified in seconds, that the advertised prefix will be valid. The default value is 86400 seconds (one day)

Default Preferred Lifetime Length of time, specified in seconds, that the client addresses generated in this prefix using SLAAC are valid. The default value is 86400 seconds (one day)

RA Subnets This section allows defining a list of subnets for which this firewall will send RA packets. Enter as many subnets as needed, each with an appropriate prefix (typically *64*.). To create an additional row for another subnet, click + Add.

DNS Settings Obtaining DNS information from RA messages is not universally supported, but for clients that do support it, using SLAAC to give an IP address and DNS from RA can do away with the need for using DHCPv6 entirely.

DNS Servers Enter up to three IP addresses for DNS Servers, or leave the fields blank to use the system default DNS servers or DNS Resolver/DNS forwarder if enabled.

Domain Search List Operates identically to the DHCP option of the same name.

Use same settings as DHCPv6 server When checked, these values will be pulled from the DHCPv6 options automatically.

21.3. IPv6 Router Advertisements

16.4 DHCPv4 & DHCPv6 Relay

DHCP requests are broadcast traffic. Broadcast traffic is limited to the broadcast domain where it is initiated. To provide DHCP service on a network segment without a DHCP server, use the DHCP relay to forward those requests to a defined server on another segment.

> Warning: It is not possible to run both a DHCP server and a DHCP Relay at the same time. To enable the DHCP relay, first disable the DHCP server on each interface.

To configure the DHCP Relay:

- Disable the DHCP Server on every interface
- Navigate to Services > DHCP Relay
- Click the tab for the interface to use with DHCP Relay
- Configure the options as follows: Enable DHCP Relay Checked

  Append circuit ID and agent ID to requests Check this to add a circuit ID (AZTCO-FW® interface number) and the agent ID to the DHCP request. This may be required by the DHCP server on the other side, and can help distinguish where the requests originated.

  Destination Server A manual entry box to set the target DHCP server

- Click Save

The DHCPv6 Relay function works identically to the DHCP Relay function for IPv4.

16.5 DNS Resolver

The DNS Resolver in AZTCO-FW® utilizes unbound, which is a validating, recursive, caching DNS resolver that supports DNSSEC and a wide variety of options. The DNS Resolver is enabled by default in current versions of AZTCO-FW.

By default, the DNS Resolver queries the root DNS servers directly and does not use DNS servers configured under System > General Setup or those obtained automatically from a dynamic WAN. This behavior may be changed, however, using the DNS Query Forwarding option. By contacting the roots directly by default, it eliminates many issues typically encountered by users with incorrect local DNS configurations, and the DNS results are more trustworthy and verifiable with Domain Name System Security Extensions (DNSSEC).

16.5.1 DNS Resolver Advanced Options

AZTCO-FW® software provides a GUI to configure some of the more common advanced options available in unbound. The options below are documented as found in the unbound.conf man page.

Hide Identity When set, attempts to query the server identity (id.server and hostname.bind) are refused.

Hide Version When set, attempts to query the server version (version.server and version.
bind) are refused.

Prefetch Support When enabled, message cache elements are prefetched before they expire to help keep the cache up to date. This option can cause an increase of around 10% more DNS traffic and load on the server, but frequently requested items will not expire from the cache.

Prefetch DNS Key Support When enabled, DNSKEYs are fetched earlier in the validation process when a Delegation Signer record is encountered. This helps lower the latency of requests but utilizes a little more CPU, and requires the cache to be set above zero.

Harden DNSSEC Data If this option is disabled and no DNSSEC data is received, then the zone is made insecure. DNSSEC data is required for trust-anchored zones. If such data is absent, the zone becomes bogus.

Message Cache Size The message cache stores DNS response codes and validation statuses. The resource record set (RRSet) cache will automatically be set to twice this amount. The RRSet cache contains the actual resource record data. The default is *4 MB*.

Outgoing TCP Buffers The number of outgoing TCP buffers to allocate per thread. The default value is *10*. If set to *0*, TCP queries will not be sent to authoritative servers.

Incoming TCP Buffers The number of incoming TCP buffers to allocate per thread. The default value is *10*. If set to *0*, TCP queries will not be accepted from clients.

EDNS Buffer Size Number of bytes size to advertise as the EDNS reassembly buffer size. This value is placed in UDP datagrams sent to peers. RFC recommendation is *4096* (the default). If fragmentation reassembly problems occur, usually observed as timeouts, then a value of *1480* may help. The *512* value bypasses most MTU path problems, but it is excessive and can generate an excessive amount of TCP fallback.

Number of Queries per Thread The number of queries that every thread will service simultaneously. If additional queries arrive that need to be serviced, and no queries can be jostled out, the new queries are dropped

Jostle Timeout Timeout used when the server is very busy. This protects against denial of service by slow queries or high query rates. The default value is *200* milliseconds. Set to a value that approximates the round-trip time to the authority servers. As new queries arrive, 50% are allowed to run and 50% are replaced by new queries if they are older than the stated timeout.

Maximum TTL for RRsets and Messages The Maximum Time to Live (TTL) for RRsets and messages in the cache, specified in seconds. The default is 86400 seconds (1 day). When the internal TTL expires the cache item is expired. This can be configured to force the resolver to query for data more often and not trust (very large) TTL values

Minimum TTL for RRsets and Messages The Minimum Time to Live for RRsets and messages in the cache, specified in seconds. The default is 0 seconds. If a record has a TTL lower than the configured minimum value, data can be cached for longer than the domain owner intended, and thus less queries are made to look up the data. The 0 value ensures the data in the cache is not kept longer than the domain owner intended. High values can lead to trouble as the data in the cache may not match up with the actual data if it changes.

TTL for Host Cache Entries Time to Live, in seconds, for entries in the infrastructure host cache. The infrastructure host cache contains round trip timing, lameness, and EDNS support information for DNS servers. The default value is *15 minutes*.

Number of Hosts to Cache Number of infrastructure hosts for which information is cached. The default is *10,000*.

Unwanted Reply Threshold If enabled, a total number of unwanted replies is tracked in every thread. When the threshold is reached, a defensive action is taken and a warning is printed to the log file. The defensive action is to clear the RRSet and message caches, hopefully flushing away any poison. The default is disabled, but if enabled a value of 10 million is suggested.

Log Level Select the log verbosity. Default is *Level 1*.

Level 0 No verbosity, only errors.

Level 1 Operational information.

Level 2 Detailed operational information.

Level 3 Query level information, output per query.

Level 4 Algorithm level information.

Level 5 Logs client identification for cache misses.

Disable Auto-added Access Control Disables the automatically-added access control entries. By default, IPv4 and IPv6 networks residing on internal interfaces of this firewall are permitted. Allowed networks must be manually configured on the Access Lists tab if when checked.

Experimental Bit 0x20 Support Use 0x20-encoded random bits in the DNS query to foil spoofing attempts. See the implementation draft dns-0x20 for more information:

16.5.2 DNS Resolver Access Lists

Unbound requires access lists (ACLs) to control which clients are allowed to submit queries. By default, IPv4 and IPv6 networks residing on internal interfaces of this firewall are permitted. Additional networks must be allowed manually.

Note: The automatic ACLs may be disabled using the Disable Auto-added Access Control option on the Advanced Settings tab.

To manage Access Lists for the DNS Resolver, navigate to Services > DNS Resolver, Access Lists tab. From this list, new entries may be added and existing entries may be edited or deleted.

When adding or editing an entry, the following options are available:

Access List Name The name for the Access List, which appears as a comment in the access list configuration file.

Action Method for handling the networks contained in this Access List

Deny Stops queries from clients in the configured networks

Refuse Stops queries from clients in the configured networks and sends back a REFUSED response code

Allow Allows queries from clients in the configured networks

Allow Snoop Allows recursive and nonrecursive queries from clients in the configured networks, used for cache snooping, and typically only configured on administrative hosts.

Description A longer text field for reference notes about this entry.

Networks A list of networks to be governed by this access list entry.

16.5.3 DNS Resolver and IPv6

The DNS Resolver is fully compatible with IPv6. It accepts and makes queries on IPv6, supports AAAA records, and has no known issues with any aspect of IPv6 and handling DNS.

16.5.4 DNS Resolver Configuration

To configure the DNS Resolver, navigate to Services > DNS Resolver

Enable Checking this box turns on the DNS Resolver, or uncheck to disable this functionality. The DNS Forwarder and DNS Resolver cannot both be active at the same time on the same port, so disable the DNS Forwarder or move one service or the other to a different port before attempting to enable the DNS Resolver.

Listen Port By default, the DNS Resolver listens on TCP and UDP port 53. This is normal for any DNS server, as it is the port clients will try to use. There are some cases where moving the DNS Resolver to another Listen Port, such as 5353 or 54 is desirable, and then specific sources may be forwarded there via port forwards.

Interfaces By default, the DNS Resolver listens on every available interface and IPv4 and IPv6 address. The Interface control limits the interfaces where the DNS forwarder will accept and answer queries. This can be used to increase security in addition to firewall rules. If a specific interface is selected, both the IPv4 and IPv6 addresses on that interface will be used for answering queries. The unbound daemon will only bind to the selected interface. Queries sent to other IP addresses on the firewall will be silently discarded.

Outgoing Network Interfaces By default the DNS Resolver utilizes all interfaces for outbound queries, so it will source the query from whichever interface and IP address is closest to the target server from a routing perspective. Selecting specific interfaces will limit the choices to only specific interfaces that may be used as a source of queries.

System Domain Local Zone Type This option determines the type of local- zone configured in unbound for the system domain. The zone type governs the type of response to give clients when there is no match in local data such as Host Overrides, DHCP hosts, etc. In each case, if there is a local match, the query is answered normally. The available types to govern non- matching responses are:

Deny Drops the query and does not answer the client.

Refuse Notifies the client that the query was refused (Using rcode REFUSED).

Static Returns a NODATA or NXDOMAIN response to the client.

Transparent This is the default behavior. If the query is for a name that does not exist locally, it is resolved as usual. If the name has a local match but the type is different, a NOERROR, NODATA response is sent to the client

Type Transparent Similar to transparent, it also passes through queries where the name matches but the type does not. For example, if a client queries for an AAAA record but only an A record exists, the AAAA query is passed on rather than receiving a negative response.

Redirect Handles queries from local data and redirects queries for zones underneath the local zone (e.g. subdomains). This can be used to control queries for all subdomains under the given domain.

Inform Answers normally, but logs the client query.

Inform Deny Denies and logs the query.

No default Disables any default content for the zone without affecting query behavior.

DNSSEC Enables Domain Name System Security Extensions (DNSSEC), which allows clients to trust the origin and content of DNS responses. This is enabled by default. DNSSEC protects against manipulation of DNS responses, such as DNS cache poisoning or other query interception, but it does not make the contents of responses secret. DNSSEC works best when using the root servers directly, unless the forwarding servers support DNSSEC. If upstream DNS servers do not support DNSSEC in forwarding mode or with domain overrides, DNS queries are known to be intercepted upstream, or clients have issues with over-size DNS responses, DNSSEC may need to be disabled.

DNS Query Forwarding Disabled by default. When enabled, unbound will use the system DNS servers from System > General Setup or those received from a dynamic WAN, rather than using the root servers directly. This is better for a multi-WAN scenario where fine control of DNS query routing is desired, but typically also requires disabling DNSSEC due to a lack of support by upstream DNS servers or other problems forwarding the queries.

DHCP Registration When active, internal machine names for DHCP clients can be resolved using DNS. This only works for clients that specify a hostname in their DHCP requests. The domain name from System > General Setup is used as the domain name on the hosts.

Static DHCP This works the same as Register DHCP leases in DNS forwarder, except that it registers the DHCP static mapping addresses instead.

Custom Options A text area for placing advanced directives for unbound that are not supported by the GUI directly. If unbound does not start correctly after entering custom options, add server: on a line before the custom options.

Host Overrides

Custom DNS entries can be created in the Host Overrides section of the page. Host overrides can define new records, or override existing records so that local clients receive the configured responses instead of responses from upstream DNS servers. This is also useful for split DNS configurations (see *Split DNS*), and as a semi-effective means of blocking access to certain specific websites.

Multiple records may be defined for the same hostname, and all IP addresses will be returned in the result. This can be used to supply both an IPv4 (A) and IPv6 (AAAA) result for a single hostname.

Note: We do not recommend using only the DNS override functionality as a means of blocking access to certain sites. There are countless ways to get around this. It will stop non-technical users, but it is easy to circumvent for those with more technical aptitude.

Host This field defines only the hostname portion of the DNS record (without the domain), e.g. www. It may be left blank to make an override record for the domain itself (Similar to an “@” record in bind.)

Domain This field is required, and defines the domain name for the override entry, e.g. example.com.

IP Address The IP address (either IPv4 or IPv6) to return as the result for a DNS lookup of this entry.

Description A text description used to identify or give more information about this entry.

Additional Names for This Host Defines additional hostnames for the same IP address (much like CNAME records) to keep them in a single override entry.

Domain Overrides

Domain overrides are found at the bottom of the DNS Resolver page. These entries specify an alternate DNS server to use for resolving a specific domain.

One example of where this is commonly deployed is in small business networks with a single internal server with Active Directory, usually Microsoft Small Business Server. The DNS requests for the Active Directory domain name must be resolved by the internal Windows Server for Active Directory to function properly. Adding an override for the Active Directory domain pointing to the internal Windows server IP address ensures these records are resolved properly whether clients are using this firewall as a DNS server or the Windows Server directly.

In an Active Directory environment the best practice is to have clients always use the Windows DNS server as the primary DNS server so dynamic name registration and other domain-related DNS tasks function properly. In environments with only one Windows DNS server, enable the DNS Resolver with an override for the Active Directory domain and use this firewall as the secondary DNS server for the internal machines. This ensures DNS resolution (except for Active Directory) does not have a single point of failure, and loss of the single server won't mean a complete Internet outage. The loss of a single server in such an environment will usually have significant consequences, but users will be more apt to leave the administrator alone to fix the problem if they can still check out their lolcats, Facebook, Twitter, et al in the meantime.

Another common use of DNS overrides is to resolve internal DNS domains at remote sites using a DNS server at the main site accessible over VPN. In such environments all DNS queries are typically resolved at the central site for centralized control over DNS, however some organizations prefer letting Internet DNS resolve with AZTCO-FW at each site, and only forwarding queries for internal domains to the central DNS server. Note a static route is necessary for this to function over IPsec. See AZTCO-FW*-initiated Traffic and IPsec* for more information.

Domain The Domain field sets the domain name that will be resolved using this entry. This does not have to be a valid TLD, it can be anything (e.g. local, test, lab), or it can be an actual domain name (example.com).

IP Address Specifies the IP Address of the DNS server to which the queries for hostnames in Domain are sent. If the target DNS server is running on a port other than 53, add the port number after the IP address with an @ separating the values, for example:

```
192.0.2.3@5353
```

Description A text description used to identify or give more information about this entry.

16.5.5 DNS Resolver and Multi-WAN

With the default settings, the DNS Resolver will have issues in a Multi-WAN environment. The main issue is that the DNS Resolver wants to query the root DNS servers directly. These queries will only be sent out using the default gateway. If the WAN containing the default gateway fails then DNS queries will also likely fail. There are ways to work around this limitation, however:

Forwarding Mode

Enable DNS Query Forwarding and configure at least one DNS server per WAN gateway under System > General Setup. DNSSEC may also need to be disabled, depending on upstream DNS server support.

Default Gateway Switching

Enable Default Gateway Switching under System > Advanced, Miscellaneous tab. This will move the default gateway to the next available gateway if the preferred default fails. However, this option is still considered experimental and may have problems in certain cases.

16.5.6 DNS Resolver and DNS Rebinding Protection

By default, DNS Rebinding protection is enabled and private IP address responses are rejected. To allow private IP address responses from a known domain, use the Custom Options box in the DNS Resolver settings to configure allowed domains as follows:

```
server: private-domain: "example.com"
```

16.5.7 CLI Commands

Unbound provides various command line utilities to manage the DNS Cache server. The following control commands are currently not available in the webGUI but can be executed from the command line.

Note: Unbound does not use the default conf file location; Use the -c flag to tell Unbound the configuration file location:

```
unbound-control -c /var/unbound/unbound.conf <unbound-command-to-run>
```

To remove items from the cache:

unbound-control flush <name> Removes <name> from the cache, all record types which include A, AAAA, NS< SOA, CNAME, DNAME, MX, PTR, SRV and NAPTR records.

unbound-control flush_type <name> <type> Removes the <name> and <type> from the cache where <type> is a particular record type.

unbound-control flush_zone <name> Removes all information at or below the <name> from cache. For example, .com will remove all entries below .com. Note this process is slow as the entire cache must be inspected.

To determine the name servers Unbound will Query to lookup a zone:

```
unbound-control lookup <name>
```

16.6 DNS Forwarder

The DNS Forwarder in AZTCO-FW® software is a caching DNS resolver that employs the dnsmasq daemon. It is disabled by default in current versions, with the *DNS Resolver* (unbound) being active by default instead. The DNS Forwarder will remain enabled on older systems or upgraded systems where it was active previously.

The DNS Forwarder uses DNS servers configured at System > General Setup, or those obtained automatically from an ISP for dynamically configured WAN interfaces (DHCP, PPPoE, PPTP). For static IP address WAN connections, DNS servers must be entered at System > General Setup or during the setup wizard for the DNS forwarder to function. Statically configured DNS servers may also be used with dynamically configured WAN interfaces by unchecking the Allow DNS server list to be overridden by DHCP/PPP on WAN box on the System > General Setup page.

By default, the DNS Forwarder queries all DNS servers at once, and the only the first response received is used and cached. This results in much faster DNS service from a client perspective, and can help smooth over problems that stem from DNS servers which are intermittently slow or have high latency, especially in Multi-WAN environments. This behavior can be disabled by activating the Query DNS servers sequentially option.

16.6.1 DNS Forwarder and IPv6

The DNS Forwarder is fully compatible with IPv6. It accepts and makes queries on IPv6, supports AAAA records, and has no known issues with any aspect of IPv6 and handling DNS.

16.6.2 DNS Forwarder Configuration

To configure the DNS Forwarder, navigate to Services > DNS Forwarder

The available options for the DNS Forwarder are:

Enable Checking this box turns on the DNS Forwarder, or uncheck to disable this functionality. The DNS Forwarder and DNS Resolver cannot both be active at the same time on the same port, so disable the DNS Resolver or move one service or the other to a different port before attempting to enable the DNS Forwarder.

DHCP Registration When active, internal machine names for DHCP clients can be resolved using DNS. This only works for clients that specify a hostname in their DHCP requests. The domain name from System > General Setup is used as the domain name on the hosts.

Static DHCP This works the same as Register DHCP leases in DNS forwarder, except that it registers the DHCP static mapping addresses instead.

Prefer DHCP When one IP address has multiple hostnames, doing a reverse lookup may give an unexpected result if one of the hostname is in host overrides and the system uses another hostname over DHCP. Checking this option will place the DHCP obtained hostnames above the static mappings in the hosts file on the firewall, causing them to be consulted first. This only affects reverse lookups (PTR), since they only return the first result and not multiple. For example, this would yield a result of labserver01.example.com, a test server’s DHCP obtained IP address, rather than a host override name of testwww.example.com that would be returned otherwise.

Query DNS servers sequentially By default, the firewall queries all DNS servers simultaneously and uses the fastest result. This isn’t always desirable, especially if there is a local DNS server with custom hostnames that could by bypassed by using a faster but public DNS server. Checking this option causes queries to be made to each DNS server in sequence from the top down, and the firewall waits for a timeout before moving on to the next DNS server in the list.

Require domain Requires a domain name on hostnames to be forwarded to upstream DNS servers. Hosts without a name will still be checked against host overrides and DHCP results, but they will not be queried against the name servers configured on the firewall. Instead, if a short hostname does not exist locally, an NXDOMAIN result ("Not Found") is returned to the client.

Do not forward private reverse lookups When checked, this option prevents dnsmasq from making reverse DNS (PTR Record) lookups for RFC1918 private IP addresses to upstream name servers. It will still return results from local entries. It is possible to use a domain override entry for the reverse lookup zone, e.g. 1.168.192 .in-addr.arpa, so that queries for a specific subnet will still be sent to a specific DNS server.

Listen Port By default, the DNS Forwarder listens on TCP and UDP port 53. This is normal for any DNS server, as it is the port clients will try to use. There are some cases where moving the DNS Forwarder to another Listen Port, such as 5353 or 54 is desirable, and then specific queries may be forwarded there via port forwards.

Interfaces By default, the DNS Forwarder listens on every available interface and all available IPv4 and IPv6 addresses. The Interface control limits the interfaces where the DNS forwarder will accept and answer queries. This can be used to increase security in addition to firewall rules. If a specific interface is selected, both the IPv4 and IPv6 addresses on that interface will be used for answering queries. Queries sent to other IP addresses on the firewall will be silently discarded.

Strict Interface Binding When set, the DNS forwarder will only bind to the interfaces containing the IP addresses selected in the Interface control, rather than binding to all interfaces and discarding queries to other addresses. This can be used similarly to the Listen Port for controlling the way that the service binds so that it can coexist with other DNS services that have similar options.

Note: This option is not compatible with IPv6 in the current version of the DNS Forwarder daemon, dnsmasq. If this is checked, the dnsmasq process will not bind to any IPv6 addresses.

Advanced Options

Custom dnsmasq configuration parameters that are not configurable in the GUI can be placed in Advanced Options. For example, to set a lower TTL for DNS records, enter max-ttl=30. Or craft a wild card DNS record to resolve .lab.example.com to 192.2.5.6 by specifying address=/lab.example.com/192.2.5.6.

Separate commands by either a space or a newline. For more information on the possible parameters that may be used, consult the dnsmasq documentation.

Host Overrides

Host override entries provide a means to configure customized DNS entries. The configuration is identical to *Host Overrides* in the DNS Resolver, refer there for details.

Domain Overrides

Domain overrides configure an alternate DNS server to use for resolving a specific domain. The configuration is identical to *Domain Overrides* in the DNS Resolver, with some slight differences:

Domain The Domain field sets the domain name that will be resolved using this entry. This does not have to be a valid TLD, it can be anything (e.g. local, test, lab), or it can be an actual domain name (example.com).

IP Address This field can be used in one of three ways. First, it can be used to specify the IP Address of the DNS server to which the queries for hostnames in Domain are sent. Second, it can be used to override another entry by entering #. For example, to forward example.com to 192.2.66. 2, but have lab.example.com forward on to the standard name servers, enter a # in this field. Third, it can be used to prevent non- local lookups by entering a !. If host override entries exist for www.example.org and mail.example.org, but other lookups for hosts under *example.org* must not be forwarded on to remote DNS servers, enter a ! in this field.

Source IP This field is optional, and primarily used to contact a DNS server across a VPN. Typically only specific local IP addresses are able to traverse a VPN, this field specifies which IP address on the firewall is used to source the DNS so the queries will pass properly.

Description A text description used to identify or give more information about this entry.

16.6.3 DNS Forwarder and Multi-WAN

The DNS Forwarder is fully compatible with Multi-WAN. Configure at least one DNS server per WAN gateway under System > General Setup.

16.6.4 DNS Forwarder and DNS Rebinding Protection

By default, DNS Rebinding protection is enabled and private IP address responses are rejected. To allow private IP address responses from a known domain, use the Advanced Options box in the DNS Forwarder settings to configure allowed domains as follows:

```
rebind-domain-ok=/example.com/
```

16.7 Dynamic DNS

The Dynamic DNS client built into AZTCO-FW® software registers the IP address of a WAN interface with a variety of dynamic DNS service providers. This is used to remotely access services on hosts that have WANs with dynamic IP addresses, most commonly VPNs, web servers, and so on.

Any number of Dynamic DNS clients may be configured using any of over 20 different Dynamic DNS providers, or even custom Dynamic DNS providers. Dynamic DNS clients can use any WAN, and can even register the real public IP address in environments where the firewall receives a private IP address for its WAN and is NATed upstream.

In addition to the typical HTTP/HTTPS-based Dynamic DNS providers, AZTCO-FW also supports RFC 2136 style Dynamic DNS updates directly to DNS servers.

16.7.1 Configuring a Dynamic DNS Client

AZTCO-FW® software allows registration with many different dynamic DNS providers. The available providers may be viewed by clicking the Service Type selector. More information about the providers may be found by searching for their name to find their web site. Several offer a basic level service at no cost, and some offer additional premium services at a cost. There is also a *Custom* option that allows for a custom URL to accommodate an unsupported provider.

Select a provider, visit their website, register for an account, and setup a hostname. The procedures for this vary with each provider, but they all have instructions on their websites. After configuring a hostname with a provider, configure AZTCO-FW with matching settings.

Most providers have the same, or similar options. There are a few types with custom options that will be covered later in this section.

To configure a Dynamic DNS client:

- Navigate to Services > Dynamic DNS
- Click ➕ Add to add a new entry
- Configure the options as follows:

  Disable Check to disable the entry, or leave unchecked so it will be active.

  Service Type Select the dynamic DNS provider here.

  Interface to Monitor Select the interface that has the IP address to keep updated, such as WAN, or an OPTx interface. Selecting a gateway group for the interface allows the Dynamic DNS entry to switch between WANs so it can allow inbound Multi-WAN failover of services on this hostname.

  Hostname Enter the hostname created at the dynamic DNS provider. This is typically the complete fully qualified domain name, such as myhost.example.com, except for Namecheap where this is only the host portion of the address.

  Domain Name For Namecheap hosts, this box must be set to the domain part of the full hostname.

  MX An MX (Mail Exchanger) record is how Internet mail servers know where to deliver mail for a domain. Some dynamic DNS providers will let MX records be configured via the dynamic DNS client. If the chosen provider allows this, enter the host name of the mail server that will receive Internet mail for the dynamic DNS domain.

  Wildcards When wildcard DNS is enabled on a dynamic DNS name, all host name queries under the given domain will resolve to the IP address of the dynamic DNS host name. For example, if the host name is example.dyndns.org, enabling wildcard will make *.example.dyndns. org (a.example.dyndns.org, b.example.dyndns.org, etc.) resolve the same as example.dyndns.org.

  Verbose Logging Check this option to increase the logging for the Dynamic DNS update process, which is useful for troubleshooting update problems.

  Verify SSL Peer When checked, the SSL certificate of the DynDNS provider server will be validated. Some servers with self-signed certificates, or those using a less common CA, may require this to be set.

  Username Enter the username for the dynamic DNS provider. Provider-specific requirements:

    Namecheap, FreeDNS Leave blank

    Route 53 Enter the Access Key ID

    GleSYS Enter the API user

Custom The username is used with basic HTTP authentication and may be left blank.

Password Enter the password for the dynamic DNS provider. Provider-specific requirements:

Namecheap, FreeDNS This is the Authentication Token

Route 53 Enter the Secret Access Key GleSYS

Enter the API Key

DNSimple Enter the API Token

Description A text field for reference.

- Click Save

Providers with Extra or Different Settings

Some providers have special settings or certain fields that need to be set in a specific way that may not be obvious. The differences are outlined in this section.

Namecheap

As mentioned earlier in the settings above, Namecheap requires that the fully qualified domain name be split into the hostname part and domain name part in separate fields.

When setting up Dynamic DNS for a *Namecheap* domain, an authentication token is given by Namecheap. This goes in the Password field, and the Username field is left blank.

HE.net Tunnelbroker

The *HE.net Tunnelbroker* choice updates an IPv6 tunnel endpoint IP address when the WAN IP changes. The Hostname in this case is the Tunnel ID from HE.net.

Route 53

When using an Amazon *Route 53* type, the Username is the Access Key ID provided by Amazon.

The following additional options are available when using *Route 53*:

Verify SSL Peer Enable to verify the server certificate when using HTTPS Zone ID

Received when creating the domain in Route 53. Must be filled in.

TTL Time to Live for the DNS record.

Custom

The *Custom* Dynamic DNS type configures options that allow for updating otherwise unsupported services. When using the custom Dynamic DNS type, the Username and Password fields are sent using HTTP basic authentication.

The following additional options are available when using *Custom*:

Interface to send update from Almost always the same as the Interface, but can be changed as needed.

Force IPv4 Resolving When checked, the update host will only be resolved using IPv4

Verify SSL Peer Enable to verify the server certificate when using HTTPS

Update URL The URL given by the Dynamic DNS provider for updates. If the IP address must appear in the URL, enter it as %IP% and the real value will be substituted as needed.

Result Match Defines expected output from the Dynamic DNS query. If it succeeds and matches the output given, then AZTCO-FW will know that the update was successful. If it does not match exactly, then it is assumed that the update failed. Leave empty to disable result checking.

DNSSimple

Verify SSL Peer Enable to verify the server certificate when using HTTPS

Zone ID Received when creating the domain.

TTL Time to Live for the DNS record.

16.7.2 Configuring RFC 2136 Dynamic DNS updates

RFC 2136 Dynamic DNS registers a hostname on any DNS server supporting RFC 2136 style updates. This can be used to update DNS records on BIND and Windows Server DNS servers, amongst others.

RFC 2136 Dynamic DNS entries may be used at the same time as regular style Dynamic DNS service providers, and like those, any number of entries can be created. RFC 2136 will update the A record, and the AAAA record if IPv6 is configured on the monitored interface.

See also:

Configuring the server infrastructure for RFC 2136 Dynamic DNS hosting is beyond the scope of this documentation, but there is a basic how- to in the recipes section: *Configuring BIND as an RFC 2136 Dynamic DNS Server*.

To configure an RFC 2136 Dynamic DNS client:

- Navigate to Services > Dynamic DNS
- Click the RFC 2136 tab
- Click ➕ Add to add a new entry
- Configure the options as follows:

  Enable Controls whether or not the entry is active. If it is unchecked, updates will not be performed for this entry.

  Interface The IP address on the chosen interface will be sent when performing the DNS update.

  Hostname The fully qualified domain name (FQDN) of the dynamic DNS entry to update. For example, myhost.example.com.

  TTL The Time To Live for the DNS entry, in seconds. Higher values will be cached longer by other name servers, so lower values are better to be sure that DNS updates are picked up in a timely manner by other servers. Usually a value between 30 and 180 seconds is reasonable, depending on how often the IP address changes.

Key Name The name of the key as specified in the DNS server configuration. For Host keys, this is typically the FQDN, so it would be identical to the value in the Hostname field. For Zone keys this would be the name of the DNS zone.

Key Type Can be one of *Zone*, *Host* or *User*. The type of key is determined by the server, so consult the server configuration or the DNS server administrator to determine the Key Type. Typically this is set to *Host*.

Key Contains the actual text of the key, e.g. /0/4bxF9A08n/zke/vANyQ==. This value is generated by the DNS server or administrator.

Server The IP address or hostname of the DNS server to which updates are sent.

Protocol When unchecked, the DNS update is sent over UDP, when checked it uses TCP instead.

Use Public IP By default, the interface IP address is always sent to the name server for the DNS update. If this box is checked, when a private IP address is detected on the selected Interface, a check is done to determine what the actual public IP address is, and then that IP address is used for the DNS update.

Record Type Determines which record(s) will be updated for this entry. For the IPv4 address, use *A*, for IPv6, use *AAAA*, or choose *Both*.

Description A free-text description of the entry for reference.

As with the other Dynamic DNS types, RFC 2136 updates are performed only when an IP address change is detected, or once every 25 days.

16.7.3 Configuring IP Address Check Services for Dynamic DNS

AZTCO-FW® software version 2.3.3 and later support custom IP address check services. These services are used by Dynamic DNS clients to determine the public IP address of the firewall when a WAN interface is behind an upstream NAT device.

To create or edit one of these services, navigate to Services > Dynamic DNS on the Check IP Services tab.

Settings

Fill out the form fields on the page as follows:

- Enable: Allow this service to be used by Dynamic DNS clients
- Name: A short name to identify this service
- URL: The full URL to the IP address check page
- Username/Password: Optional authentication to use when accessing the URL
- Verify SSL Peer: Check this box if the server has a self-signed SSL certificate or a certificate from a CA that is not trusted by the firewall.
- Description: A longer description of this service

Once a service is defined, it may be selected on individual Dynamic DNS service entries.

Server-Side Configuration Examples

Hosting one of these services is very simple. The server page need only print the requesting client IP address in the expected format:

```
Current IP Address: x.x.x.x
```

nginx (internal/native)

```
location /ip { default_type text/html; return 200 "<html><head><title>Current IP Check</title></head><body>Current IP
↪Address: $remote_addr</body></html>";
}
```

nginx (internal with LUA)

```
location = /ip { default_type text/html; content_by_lua ' ngx.say("<html><head><title>Current IP
        Check</title></head><body>Current IP␣
 ↪Address: ") ngx.say(ngx.var.remote_addr)
                ngx.say("</body></html>")
        ';
}
```

PHP

```
<html><head><title>Current IP Check</title></head><body>Current IP Address: <?=$_
↪SERVER['REMOTE_ADDR']?></body></html>
```

16.8 SNMP

The Simple Network Management Protocol (SNMP) daemon enables remote monitoring of some AZTCO-FW® system parameters. Depending on the options chosen, monitoring may be performed for network traffic, network flows, pf queues, and general system information such as CPU, memory, and disk usage. The SNMP implementation used by AZTCO-FW is bsnmpd, which by default only has the most basic management information bases (MIBs) available, and is extended by loadable modules. In addition to acting as an SNMP daemon, it can also send traps to an SNMP server for certain events. These vary based on the modules loaded. For example, network link state changes will generate a trap if the MIB II module is loaded.

The SNMP service can be configured by navigating to Services > SNMP.

The easiest way to see the available data is to run snmpwalk against the firewall from another host with net-snmp or an equivalent package installed. The full contents of the MIBs available are beyond the scope of this documentation, but there are plenty of print and online resources for SNMP, and some of the MIB trees are covered in RFCs. For example, the Host Resources MIB is defined by RFC 2790.

21.8.1 SNMP and IPv6

The bsnmpd daemon does not currently support IPv6.

16.8.2 SNMP Daemon

These options dictate if, and how, the SNMP daemon will run. To turn the SNMP daemon on, check Enable. Once Enable has been checked, the other options may then be changed.

Polling Port SNMP connections are made using only UDP, and SNMP clients default to using UDP port 161. This setting controls which port is used for the SNMP daemon, and the SNMP client or polling agent must be changed to match.

System location This text field specifies a string to return when the system location is queried via SNMP. Any text may be used here. For some devices a city or state may be close enough, while others may need more specific detail such as which rack and position in which the system resides.

System contact A string defining contact information for the system. It can be a name, an e-mail address, a phone number, or whatever is needed.

Read Community String With SNMP, the community string acts as a kind of username and password in one. SNMP clients will need to use this community string when polling. The default value of public is common, so we strongly recommend using a different value in addition to restricting access to the SNMP service with firewall rules.

16.8.3 SNMP Traps

To instruct the SNMP daemon to send SNMP traps, check Enable. Once Enable has been checked, the other options may then be changed.

Trap server The trap server is the hostname or IP address to which SNMP traps are forwarded.

Trap server port By default, SNMP traps are set on UDP port 162. If the SNMP trap receiver is set for a different port, adjust this setting to match.

SNMP trap string This string will be sent along with any SNMP trap that is generated.

16.8.4 Modules

Loadable modules allow the SNMP daemon to understand and respond to queries for more system information. Each loaded module will consume additional resources. As such, ensure that only required modules are loaded.

MibII This module provides information specified in the standard MIB II tree, which covers networking information and interfaces. Having this module loaded will, among other things, provides network interface information including status, hardware and IP addresses, the amount of data transmitted and received, and much more.

Netgraph The netgraph module provides some netgraph-related information such as netgraph node names and statuses, hook peers, and errors.

PF The pf module provides a wealth of information about pf. The MIB tree covers aspects of the ruleset, states, interfaces, tables, and ALTQ queues.

Host Resources This module provides information about the host itself, including uptime, load average and processes, storage types and usage, attached system devices, and even installed software. This module requires MibII, so if MibII is unchecked when this option is checked, MibII will be checked automatically.

UCD This module provides various system information knows as the ucdavis MIB, or UCD-SNMP-MIB. It provides information about memory usage, disk usage, running programs, and more.

Regex The Regex module is reserved for future use or use by users customizing the code to their needs. It allows creating SNMP counters from log files or other text files.

16.8.5 Interface Binding

This option configures the SNMP daemon to listen only on the chosen interface or virtual IP address. All interfaces with IP addresses, CARP VIPs, and IP Alias VIPs are displayed in the drop-down list.

Binding to a specific local interface can ease communication over VPN tunnels, as it eliminates the need for the previously mentioned static route, and it also provides extra security by not exposing the service to other interfaces. It can also improve communication over multiple local interfaces, since the SNMP daemon will reply from the "closest" address to a source IP address and not the IP address to which the query was sent.

16.9 UPnP & NAT-PMP

Universal Plug and Play (UPnP) and NAT Port Mapping Protocol (NAT-PMP) are network services which allow software and devices to configure each other when attaching to a network. This includes automatically creating their own dynamic NAT port forwards and associated firewall rules.

The UPnP and NAT-PMP service on AZTCO-FW®, found at Services > UPnP & NAT-PMP, enables client PCs and other devices such as game consoles to automatically allow required inbound traffic. There are many popular programs and systems which support UPnP, such as Skype, uTorrent, mIRC, IM clients, Wii U, PlayStation 4, and XBox One. NATPMP is supported on Apple products.

UPnP employs the Simple Service Discovery Protocol (SSDP) for network discovery, which uses UDP port 1900. The UPnP daemon used by AZTCO-FW, miniupnpd , also uses TCP port 2189. When using a strict LAN ruleset, manually add firewall rules to allow access to these services, especially if the default LAN-to-any rule has been removed, or in bridged configurations. NAT-PMP is also handled by miniupnpd and uses UDP port 5351.

16.9.1 UPnP & NAT-PMP and IPv6

As of this writing, the UPnP and NAT-PMP service on current versions of AZTCO-FW supports IPv6, but client support is still spotty.

16.9.2 Security Concerns

UPnP and NAT-PMP are a classic example of the "Security vs. Convenience" trade- off. By their very nature, these services are insecure. Any program on the network can allow in and forward any traffic – a potential security nightmare. On the other side, it can be a chore to enter and maintain NAT port forwards and their associated rules, especially when it comes to game consoles. There is a lot of guesswork and research involved to find the proper ports and settings, but UPnP *just works* and requires little administrative effort. Manual port forwards to accommodate these scenarios tend to be overly permissive, potentially exposing services that should not be open from the Internet. The port forwards are also always on, where UPnP may be temporary.

Access controls exist in the UPnP service configuration, which helps to lock down which devices are allowed to make alterations. Over and above the built-in access controls, further control may be exerted with firewall rules. When properly controlled, UPnP can also be a little more secure by allowing programs to pick and listen on random ports, instead of always having the same port open and forwarded.

21.9. UPnP & NAT-PMP

16.9.3 Configuration

To configure UPnP and NAT-PMP:

- Navigate to Services > UPnP & NAT-PMP
- Configure the options as follows:

Enable UPnP & NAT-PMP Master control for the entire service. When unchecked, all of the services on this page are disabled.

Allow UPnP Port Mapping When checked, UPnP is allowed.

Allow NAT-PMP Port Mapping When checked, NAT-PMP is allowed.

External Interface The WAN interface for outgoing traffic. This must be set to the WAN containing the default gateway. Only one External Interface may be selected.

Interfaces The local interfaces where clients allowed to use UPnP/NAT-PMP reside. When a bridge is in use, only select the bridge interface with an IP address. Multiple interfaces may be selected.

Download Speed Maximum download speed reported to clients, in Kilobits per second.

Upload Speed Maximum upload speed reported to clients, in Kilobits per second.

Override WAN Address Selects an alternate interface IP address to use, such as a CARP or IP Alias Virtual IP address.

Traffic Shaping Queue The name of an ALTQ (not Limiter) traffic shaping queue in which traffic allowed through using UPnP will be placed.

> Note: Exercise caution when selecting this queue. UPnP is used by traffic such as game consoles, which need high priority, and also by file transfer clients which may need low priority.

Log Packets When checked, port forwards generated by UPnP/NAT-PMP will be set to log, so that each connection made will have an entry in the firewall logs, found at Status > System Logs, on the Firewall tab.

Use System Uptime By default, the UPnP daemon reports the service uptime when queried rather than the system uptime. Checking this option will cause it to report the actual system uptime instead.

Deny Access by Default When checked, UPnP will only allow access to clients matching the access rules. This is a more secure method of controlling the service, but as discussed above, is also less convenient.

User Specified Permissions These fields specify user-defined access rules. If the default-deny option is chosen, rules must be set to allow access. Additional rules may be added by clicking ✚ Add Rules are formulated using the following format:

```
<[allow|deny]> <[external port|port range]> <[internal IP|IP/CIDR]> <[internal␣
↪port|port range]>
```

- Click Save

The UPnP and/or NAT-PMP service will be started automatically.

UPnP User Permission Examples

Deny access to external port 80 forwarding from everything on the LAN, 192.168.1.1, with a /24 subnet, to local port 80:

```
deny 80 192.168.1.1/24 80
```

Allow 192.168.1.10 to forward any unprivileged port:

```
allow 1024-65535 192.168.1.10 1024-65535
```

16.9.4 Status

The status of the UPnP daemon process may be viewed at Status > Services. The Service Status page shows if the daemon is running or stopped, and allows the service to be stopped, started or restarted. Under normal circumstances, manually managing the daemon is not necessary.

A list of currently forwarded ports and clients, similar to Figure *UPnP & NAT-PMP Status Screen Showing Client PCs With Forwarded Ports*, may be viewed under Status > UPnP & NAT-PMP.

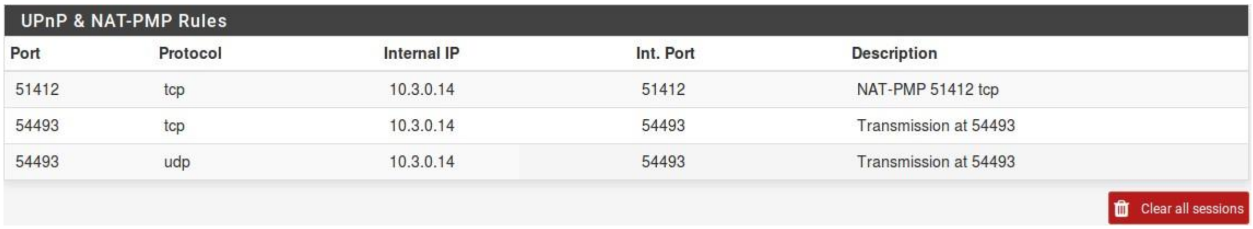

UPnP & NAT-PMP Rules

Port	Protocol	Internal IP	Int. Port	Description
51412	tcp	10.3.0.14	51412	NAT-PMP 51412 tcp
54493	tcp	10.3.0.14	54493	Transmission at 54493
54493	udp	10.3.0.14	54493	Transmission at 54493

Fig. 1: UPnP & NAT-PMP Status Screen Showing Client PCs With Forwarded Ports

16.9.5 Troubleshooting

Most issues with UPnP tend to involve bridging. In this case it is important to have firewall rules allow UPnP on UDP port 1900. Since it is multicast traffic, the destination will be the broadcast address for the subnet, or in some cases making it *any* will be necessary. Consult the firewall logs at Status > System Logs, on the Firewall tab to see if traffic is being blocked. Pay particular attention to the destination address, as it may be different than expected.

Further trouble with game consoles may also be alleviated by switching to manual outbound NAT and enabling Static Port. See *Static Port* for more details.

16.10 NTPD

The NTP service is a Network Time Protocol (NTP) daemon which will listen for requests from clients and allow them to synchronize their clock with that of the AZTCO-FW® firewall. By running a local NTP server and using it for local clients, it reduces the load on the lower-stratum servers and can ensure that local systems can always reach a time server. Before delegating this task to a firewall running AZTCO-FW, the best practice is to ensure that the firewall has an accurate clock and keeps time reasonably.

16.10.1 NTP and IPv6

The NTP Project daemon fully supports IPv6 as a client and a server.

NTP Server Configuration

To configure the NTP Server:

- Navigate to Services > NTP
- Configure the settings as follows:

  Interface Select the interface(s) to use for NTP. The NTP daemon binds to all interfaces by default to receive replies properly. This may be minimized by selecting at least one interface to bind, but that interface will also be used to source the NTP queries sent out to remote servers, not only to serve clients. Deselecting all interfaces is the equivalent of selecting all interfaces.

  Time Servers A list of servers to query in order to keep the clock of this firewall synchronized. This list is initially pulled from the entries under System > General Setup. For best results, we recommend using at least three servers, but no more than five. Click ➕ Add to configured additional time servers.

Prefer When checked, this NTP server entry is favored by the NTP daemon over others.

No Select When checked, this NTP server is not used for time synchronization, but only to display statistics.

Orphan Mode Orphan mode uses the system clock when no other clocks are available, otherwise clients will not receive a response when other servers are unreachable. The value entered here is the stratum used for Orphan Mode, and is typically set high enough that live servers are preferred. The default value is 12.

NTP Graphs Check to enable RRD graphs for NTP server statistics.

Logging When logging options are active, NTP logs are written using syslog and may be found under Status > System Logs, on the NTP tab.

Log Peer Messages When checked, NTP will log messages about peer events, information, and status.

Log System Messages When checked, NTP will log messages about system events, information, and status.

Statistics Logging Click ⚙ Show Advanced to view these options. When enabled, NTP will create persistent daily log files in /var/log/ntp to keep statistics data. The format of the statistics records in the log files can be found in the ntp.conf man page

Log reference clock statistics When checked, NTP records clock driver statistics on each update.

Log clock discipline statistics When checked, NTP records loop filter statistics on each update of the local clock.

Log NTP Peer Statistics When checked, NTP records statistics for all peers of the NTP daemon, along with special signals.

Leap Seconds Click ⚙ Show Advanced to view these options. Defines the contents of the Leap Second file, used by NTP to announce upcoming leap seconds to clients. This is typically used only by stratum 1 servers. The exact format of the file may be found on the IETF leap second list

- Click Save

Access Restrictions

Access restrictions (ACLs) are configured on the ACL tab under Services > NTP. These ACLs control how NTP interacts with clients.

Default Access Restrictions Control behavior for all clients by default.

Kiss-o'-Death When set, NTP will send a KoD packet when an access violation occurs. Such packets are rate limited and no more than one per second will be sent.

Modifications When set, ntpq and ntpdc queries that attempt to change the configuration of the server are denied, but informational queries are returned.

Queries When set, all queries from ntpq and ntpdc are denied.

> Warning: Setting this will effectively disable the NTP status page, which relies on ntpq.

Service When set, NTP will deny all packets except queries from ntpq and ntpdc.

Peer Association When set, NTP denies packets that would result in a new peer association, including broadcast and symmetric active packets for peers without an existing association.

Trap Service When set, NTP will not provide mode 6 control message trap service, used for remote event logging.

Custom Access Restrictions Defines the behavior for specific client addresses or subnets. Click ➕ Add to add a new network definition.

Network/mask The subnet and mask to define the client controlled by the restrictions in this entry.

Restrictions The option names are abbreviated versions of those in the default list, in the same order.

Click Save to store the ACLs.

Serial GPS

If this firewall has an available serial port, a Serial GPS may be used to provide a reference clock for the firewall. If the GPS also supports a Pulse Per Second (PPS) signal, that may also be used as a PPS clock reference.

> Warning: USB GPS units may function, but we do not recommend their use due to USB timing issues. The overhead of USB makes its unreliable as a clock or timing source.

For best results, we recommend configuring at least two NTP servers under System > General Setup or Services > NTP to avoid loss of sync if the GPS data is not valid over time. Otherwise the NTP daemon may only use values from the unsynchronized local clock when providing time to clients.

To configure a GPS for use by NTP:

- Navigate to Services > NTP
- Click the Serial GPS tab
- Configure the settings as follows:

  GPS Type Select the make and model of the GPS unit. If the model is unknown, use the *Default* choice. If the model is known but not listed, use *Custom*.

Serial Port All serial ports detected on the firewall are listed. Select the port with the GPS attached. Onboard hardware serial ports start with cuau, USB serial ports are prefixed with cuaU.

Baud Rate Enter the serial speed for the GPS, typically a low value such as 4800

NMEA Sentences By default, NTP will listen for all supported NMEA sentences. To limit this to specific types, select them from the list.

Fudge Time 1 Specifies a constant to be added to the GPS PPS signal as an offset.

Fudge Time 2 Specifies a constant to be added to the GPS time as an offset.

Stratum Used to configure the stratum of the GPS clock. The default value is 0 so the GPS is preferred over all others. If another clock must be preferred instead, set the stratum value higher than the stratum of the preferred clock.

Flags These options provide additional tweaks to fine-tune the GPS behavior:

Prefer this clock Marks the reference clock as preferred by NTP.

Do not use this clock Prevents the clock from being used by NTP for time synchronization, it is only displayed for reference.

PPS signal processing Enables processing of the Pulse Per Second (PPS) signal in the GPS driver. Only enable this if the GPS is known to output a usable PPS signal.

Falling edge PPS signal processing When set, the falling edge of the PPS signal is used for timing, rather than the rising edge.

Kernel PPS clock discipline When set, the OS Kernel will use PPS directly for timing.

Obscure location in timestamp Obscures the GPS data so the location of the clock cannot be determined.

Log the sub-second fraction of the received time stamp When checked,this can rapidly fill the log, but can be useful for fine tuning of Fudge Time 2.

Clock ID A 1-4 character identifier used to change the GPS Clock ID. The default value is GPS.

GPS Initialization Contains the initialization string sent to the GPS at start up to configure its behavior. When using the *Custom* GPS type, a proper initialization string for the GPS must be entered manually.

NMEA Checksum Calculator Calculates a checksum for use when crafting new GPS Initialization values or adjusting existing values.

- Click Save

PPS Source (Non-GPS)

A non-GPS PPS Source, such as a radio, may also be used for clock timing. It cannot be used for synchronization since there is no time data, but it can be used to ensure a clock ticks accurately.

To configure a Non-GPS PPS source:

- Navigate to Services > NTP
- Click the PPS tab
- Configure the settings as follows:

Serial Port All serial ports detected on the firewall are listed. Select the port with the GPS attached. Onboard hardware serial ports start with cuau, USB serial ports are prefixed with cuaU.

Fudge Time 1 Specifies a constant to be added to the PPS signal as an offset, to account for delay between the transmitter and receiver.

Stratum Used to configure the stratum of the PPS source. The default value is 0 so the PPS source is preferred over all others. If another clock must be preferred instead, set the stratum value higher than the stratum of the preferred clock.

Flags

Falling edge PPS signal processing When set, the falling edge of the PPS signal is used for timing, rather than the rising edge.

Kernel PPS clock discipline When set, the OS Kernel will use PPS directly for timing.

Record a timestamp Record a timestamp once for each second, which is useful for constructing Allan deviation plots.

Clock ID A 1-4 character identifier used to change the PPS Clock ID. The default value is PPS. •

Click Save

See also:

- *Status*

16.11 Wake on LAN

The Wake on LAN (WOL) page at Services > Wake on LAN can wake up computers from a powered-off state by sending special "Magic Packets".

The network interface card in the client computer that is to be woken up must support WOL and it must be configured properly. Typically there is a BIOS setting to enable WOL, and non-integrated adapters often require a WOL cable connected between the NIC and a WOL header on the motherboard.

WOL has many potential uses. Typically, workstations and servers are kept running because of services they provide, files or printers they share, or for convenience. Using WOL would allow these to remain in a sleep state to conserve power. When a service is required, the system can be woken up when needed. Another example would be if someone needs remote access to a system, but the user shut it down before leaving the office. Using WOL the target system can be awoken, and it may then be accessed once it has booted.

Warning: WOL offers no inherent security. Any system on the same layer 2 network may transmit a WOL packet, and the packet will be accepted and obeyed. It is best to only configure WOL in the BIOS for machines that need it, and disable it in all others. There are some vendor-specific WOL extensions that provide extra security, but nothing universally supported.

16.11.1 Wake Up a Single Machine

To wake up a single machine:

- Navigate to Services > Wake on LAN
- Select the Interface through which the target system can be reached
- Enter the target system MAC address in the format of xx:xx:xx:xx:xx:xx
- Click Send

AZTCO-FW® software will transmit a WOL Magic Packet out the chosen interface, and if everything went as planned, the system will power on and start to boot. Keep in mind that systems will take some time to boot. It may be several minutes before the target system is available.

16.11.2 Storing MAC Addresses

To store a MAC address for convenience:

- Navigate to Services > Wake on LAN
- Click ➕ Add under the list of stored MAC addresses to add a new entry
- Select the Interface through which the target system can be reached
- Enter the target system MAC address in the format of xx:xx:xx:xx:xx:xx
- Enter a Description for the entry, such as the target system's name, owner, or location. For example: "Pat's PC" or "Sue's Server"
- Click Save

Once saved, the entry will be available on the list at Services > Wake on LAN.

Maintaining the entries is similar to other tasks in AZTCO-FW: Click ⧉ to edit an existing entry, and click 🗑 to remove an entry.

16.11.3 Wake a Single Stored Machine

To send a WOL Magic Packet to a system that has been previously stored:

- Navigate to Services > Wake on LAN
- Locate the desired entry in the list
- Click its MAC address or click the ⏻ icon in the Actions column

The WOL page will reload, and the Magic Packet will be sent. The status of the WOL attempt will also be displayed.

16.11.4 Wake All Stored Machines

To send a WOL Magic Packet to all stored systems at once:

- Navigate to Services > Wake on LAN
- Click ⏻ Wake All Devices under the list of stored addresses.

16.11.5 Wake from DHCP Leases View

To send a WOL Magic Packet from the DHCP Leases view:

- Navigate to Status > DHCP Leases
- Locate the desired system in the list
- Click ⏻ at the end of the lease row to send a WOL Magic Packet

Note: The WOL function is only available for systems marked offline, meaning they are not in the ARP table on the firewall. If a system was very recently powered off, it can take a few minutes for the ARP entry to expire before it will be marked offline.

If a system has been powered off for quite some time, clicking ⊕ Show all configured leases might be required to see the previous lease.

When the link is clicked, the browser will return to the WOL page, and the Magic Packet will be sent.

16.11.6 Save from DHCP Leases View

A MAC address and hostname may be copied to a new WOL mapping entry while viewing the DHCP leases.

- Navigate to Status > DHCP Leases
- Locate the desired system in the list

- Click ⊞ at the end of lease entry
- Confirm the values on the page, and enter any missing information.
- Click Save

16.12 PPPoE Server

AZTCO-FW® software can act as a PPPoE server, accepting and authenticating connections from PPPoE clients on a local interface, in the role of an access concentrator (LAC). This feature can be used to force users to authenticate before gaining network access, or otherwise control their login behavior.

The PPPoE Server is located at Services > PPPoE Server. The configuration is very similar to the L2TP VPN server (*L2TP VPN*).

Multiple PPPoE servers may be configured on separate interfaces. To begin setting up a PPPoE server:

- Navigate to Services > PPPoE Server
- Click ✚ Add to add a new server entry
- Configure the PPPoE Server as follows:

  Enable When checked, this PPPoE Server instance will be active.

  Interface The single interface upon which PPPoE service will be available.

  Total User Count Determines how many clients in total are allowed to connect to this instance.

  User Max Logins Determines how many times a single client may login concurrently.

  Server Address The IP address which the AZTCO-FW system will send to the PPPoE clients to use as their gateway.

  > Warning: This IP address must not be an IP address currently in use on the firewall.

  Remote Address Range The IP address for the start of the PPPoE client subnet. Together with the Subnet Mask it defines the network used by the PPPoE clients.

  Subnet Mask Defines the CIDR mask assigned to PPPoE clients.

  Description Optional explanatory text for this server instance.

  DNS Servers Optional fields used to send specific DNS servers to the PPPoE clients, otherwise the firewall IP address will be sent to the client for DNS if the DNS Forwarder or DNS Resolver are enabled. If the DNS Forwarder and DNS Resolver are both disabled, then the DNS servers configured on the firewall will be sent instead.

- Configure RADIUS if that will be utilized for user authentication. Any RADIUS server may be used.

  See also:

  See *Authenticating from Active Directory using RADIUS/NPS* for information on setting up RADIUS on a Windows server.

Use RADIUS Authentication Check to configure the PPPoE server to use at least one RADIUS server for Authentication instead of local users.

Use RADIUS Accounting Optional, sends RADIUS accounting data to the RADIUS server to note items such as login and logout times, and bandwidth used.

Use a Backup RADIUS Authentication Server A second RADIUS server to use if the primary RADIUS server fails.

21.12. PPPoE Server

NAS IP Address Optional, sends a specific IP address to the RADIUS server for the NAS-IP-Address attribute.

RADIUS Accounting Update The interval at which accounting data is sent to the RADIUS server, in seconds.

RADIUS Issued IP Addresses When checked, IP addresses can be assigned to users via RADIUS reply attributes.

Primary RADIUS Server The preferred RADIUS server to use for Authentication.

IP Address The IP address of the RADIUS server

Authentication Port The port used for authentication (typically 1812)

Accounting Port The port used for accounting data (typically 1813)

Primary RADIUS Server Shared Secret The shared secret configured for this firewall on the RADIUS server. The same value must be entered in the Confirm box.

Secondary RADIUS Server Same type of settings as the primary, but defines the secondary RADIUS server.

- Add users to the server to utilize local authentication:
  - Click + Add User

    Username The username for the user account

    Password The password for the user account

    IP Address An optional static IP address to assign the user at login
  - Repeat as needed
- Click Save

16.13 IGMP Proxy

The Internet Group Management Protocol (IGMP) Proxy provides a means to proxy multicast traffic between network segments.

The IGMP Proxy service can be found at Services > IGMP Proxy.

For a working IGMP Proxy configuration, one upstream and at least one downstream interface must be defined.

To configure the IGMP Proxy:

- Navigate to Services > IGMP Proxy
- Click + Add to create a new interface instance
- Configure the instance as follows:

  Interface The interface to be used for this instance

  Description Optional text to describe this instance

  Type The type of network interface defined by this instance

    Upstream Interface The outgoing interface which is responsible for communicating to available multicast data sources. There can only be one upstream interface.

    Downstream Interface The distribution interfaces to the destination networks, where multicast clients can join groups and receive multicast data. One or more downstream interfaces must be configured.

  Threshold The TTL threshold for forwarded data on an interface, to prevent looping from occurring. Packets with a TTL lower than the value in this field will be ignored. The default TTL is 1 if the field is left blank.

  Networks A list of CIDR-masked Network entries to control what subnets are allowed to have their multicast data proxied. Click + Add Network to enter additional networks.

  – Click Save

A firewall rule is also required on the Downstream side (e.g. *LAN*) to match and pass the multicast traffic. In the *Advanced Options* of the firewall rule, Allow packets with IP Options must be enabled.

The base install of AZTCO-FW® software includes services which add fundamental functionality and flexibility to the firewall. The topics in this chapter discuss services in the base installation that the firewall provides for other hosts on the network. These services include allocating IPv4 and IPv6 addresses via DHCP, DNS resolution and Dynamic DNS, SNMP, UPnP and more. Additional services can also be added with packages.

CHAPTER

SEVENTEEN

DHCP

Dynamic Host Configuration Protocol (DHCP), allows a device such as AZTCO-FW® software to dynamically allocate IP addresses to clients from a predefined pool of addresses. DHCP also sends configuration information to clients such as a gateway, DNS servers, domain name, and other useful settings.

17.1 Using DHCP Search Domains on Windows DHCP Clients

The DNS Search Domain functionality present in the DHCP Server settings in AZTCO-FW® software is only supported by some DHCP clients; AZTCO-FW software uses the standard DHCP option 119 mechanism to deliver the search domains to clients which request them.

Unfortunately, The Microsoft Windows DHCP client does not support requesting option 119, so no matter which DHCP server is used, clients running Microsoft Windows can never receive or use a search domain list from DHCP.

If the settings must be used by clients, they can be pushed via GPO or in the extreme case, the clients can be replaced by ones which support option 119 such as BSD, Linux, OSX, and so on.

Sources:

- http://social.technet.microsoft.com/Forums/en-US/winserverNIS/thread/9ba77f86-4708-42ca-a193-2a01b813ec27/
- http://social.technet.microsoft.com/Forums/en-US/winserverNIS/thread/7ba59619-3484-43fa-8585-a2d69ccd00df/
- http://technet.microsoft.com/en-us/library/dd572752%28v=office.13%29.aspx (See comments)
- http://serverfault.com/questions/37417/which-dhcp-client-os-support-dhcp-option-119-domain-suffix-search

17.2 Static Mappings Inside DHCP Pools

While ISC dhcpd will allow a static mapping to be defined inside the DHCP range/pool, it can result in unexpected behavior.

ISC dhcpd only checks via ping to ensure that an IP is not actively in use when making assignments. Making a static mapping does not “reserve” that IP out of the pool. The static mapping in this case merely represents a preference for an IP and others are not prevented from taking the IP if it is not in use.

An example: If the DHCP pool is from 192.168.0.10 to 192.168.0.250, and a static mapping is defined for 192.168.0.25. If the PC that normally has 192.168.0.25 is ever offline another device could be assigned 192.168.0.25. When the other machine powers back up it will not be able to get 192.168.0.25 because it is currently in use.

As such, it is best to only make assignments outside the range/pool, and the AZTCO-FW® webGUI enforces this practice.

If assignments absolutely must be made inside the pool, and the risks involved are worth taking and want to do so anyway, the input validation check may be removed from the PHP file that drives the DHCP editor page. The details of this unsupported change are left out as an exercise for the reader.

See also:

- *Status*
- *Viewing DHCPv6 Leases*
- *DHCP Logs*
- *Troubleshooting DHCPv6 Client XID Mismatches*
- *Troubleshooting Offline DHCP Leases*

CHAPTER

EIGHTEEN

DNS

DNS, or Domain Name System, is the mechanism by which a network device resolves a name like www.example.com to an IP address such as 198.51.100.25, or vice versa. Clients must have functional DNS if they are to reach other devices such as servers using their hostnames or fully qualified domain names.

18.1 DNS Resolver/Forwarder

These topics cover using AZTCO-FW® software as a caching DNS resolver or forwarder, which handles DNS requests from local clients. When acting as a resolver or forwarder, AZTCO-FW software will performs DNS resolution or hand off queries to an upstream DNS forwarding server.

18.1.1 DNS Rebinding Protections

AZTCO-FW® software includes some built in methods of protection against DNS rebinding attacks. These measures are described below.

DNS forwarder

The DNS forwarder (dnsmasq) uses the option –stop-dns-rebind by default, which rejects and logs addresses from upstream nameservers which are in the private IP ranges. In the most common usage, this is filtering DNS responses received from the Internet to prevent DNS rebinding attacks. Internet DNS responses should never come back with a private IP, hence it's safest to block this.

There are some cases when public DNS servers have private IP address replies by default, though it is not recommended. In those cases, DNS rebinding can be disabled or an override may be placed in the DNS Forwarder Advanced Settings box as follows:

```
rebind-domain-ok=/mydomain.com/
```

Note this is automatically overridden for domains in the DNS forwarder's domain override list, as the most common usage of that functionality is to resolve internal DNS hostnames.

DNS Resolver (Unbound)

Unbound has similar protections to dnsmasq, using its “Private Address support” option. With that option enabled RFC1918 addresses are stripped away from DNS answers. Additionally, the DNSSEC validator may mark the answers bogus.

In the package on 2.1 and earlier this option is located in the main “Unbound DNS Settings” tab. On 2.2 where Unbound is integrated into the base system, it is active by default and controlled by the DNS Rebinding option under System > Advanced.

Individual domains can be excluded from DNS rebinding protection using the Custom Options on the Unbound general settings. Enter one domain per line in the following format, preceded by the “server:” line.

```
server: private-domain:
"example.com"
```

Web interface protection

For those not using the DNS forwarder, and as an additional layer of checks, the web interface will block attempts to access it via an unknown hostname. It will display “Potential DNS Rebind Attack Detected” and drop any request. By default, only the hostname and domain configured under System>General Setup are accepted. For instance if firewall.example.com is configured as the system’s hostname, and it is loaded in a browser using fw1.example.com, that attempt will be rejected. Additional hostnames can be added under System>Advanced, “Alternate Hostnames”.

Logging in using the IP address of the system rather than the hostname does work if this message is encountered when attempting to load by hostname. Once access has been obtained, configure the hostname(s) accordingly and then it is possible to log in using the desired hostname.

If this message is encountered when a client attempted to access a forwarded service (Port forward, 1:1 NAT, relayd, etc) it indicates that the request did not match any NAT rules. From the inside of the network, this would require NAT reflection or split DNS to accomplish.

See also:

Accessing Port Forwards from Local Networks

18.1.2 Creating Wildcard Records in DNS Forwarder/Resolver

A wildcard DNS record resolves anything.example.com to a single IP, which can be useful in certain cases.

DNS Forwarder (dnsmasq)

A wildcard entry may be created in the DNS Forwarder by using the advanced options box, and an entry like so:

```
address=/example.com/192.168.1.54
```

That would make any host under *example.com* resolve to 192.168.1.54 (www.example.com, thissitedoesnotexist.example.com, mystuff.example.com, and so on).

If a specific Host Overrides is set for example:

```
specific.example.com 192.168.1.100 knownhost.example.com
192.168.1.101
```

Then those would be returned when doing a query for those hosts, only when no specific host has been specified in the host overrides would the advanced wildcard entry be used.

To resolve the domain to an IP address:

```
example.com 192.168.1.45
```

Leave the host field blank in the host overrides. So if the query is now for *example.com*, 192.168.1.45 would be returned. If *knownhost.example.com* was queried for then 192.168.1.101 would be returned.

If a blank hostname *example.com* host override entry has not been created, then a query for *example.com* would return the wildcard IP address set in the advanced option.

If *madeupname.example.com* was queried, then since no specific host record for *madeupname* exists in the host overrides. The wildcard entry of 192.168.1.54 would be returned.

DNS Resolver (Unbound)

The same effect may be obtained in the DNS Resolver (Unbound) using its advanced options:

```
server:
local-zone: "example.com" redirect
local-data: "example.com 86400 IN A 192.168.1.54"
```

See also:

- *DNS Lookup*
- *Troubleshooting the DNS Forwarder*

18.2 DNS Guides

How to perform various tasks related to DNS.

- *Blocking External Client DNS Queries*
- *Redirecting Client DNS Requests*
- *Troubleshooting the DNS Cache*

18.3 Dynamic DNS

Dynamic DNS updates an external DNS server with an interface IP address when it changes. This enables a firewall with a dynamic WAN such as DHCP or PPPoE to host public services even when its IP address changes periodically.

CHAPTER

NINETEEN

TRAFFIC SHAPER

19.1 What the Traffic Shaper can do for a Network

The basic idea of traffic shaping is raising and lowering the priorities of packets or keeping them under a certain speed. This concept seems simple, however, the number of ways in which this concept can be applied is vast. These are but a few common examples that have proven popular with users of AZTCO-FW® software.

19.1.1 Keep Browsing Smooth

Asymmetric links, where the download speed differs from the upload speed, are commonplace, especially with DSL. Some links are so out of balance that the maximum download speed is almost unattainable because it is difficult for a firewall to send out enough ACK (acknowledgement) packets to keep traffic flowing. ACK packets are transmitted back to the sender by the receiving host to indicate that data was successfully received, and to signal that it is OK to send more. If the sender does not receive ACKs in a timely manner, congestion control mechanisms in TCP will kick in and slow down the connection.

This type of situation is common: When uploading a file over a link that has asymmetric throughput capability, browsing and downloading slows to a crawl or stalls. This happens because the uploading portion of the circuit is full from the file upload and there is little room to send ACK packets which allow downloads keep flowing. By using the shaper to prioritize ACK packets, the firewall can enable faster, more stable download speeds on asymmetric links.

This is not as important on symmetric links where the upload and download speed are the same, but may still be desirable if the available outgoing bandwidth is heavily utilized.

19.1.2 Keep VoIP Calls Clear

If Voice over IP calls use the same circuit as data, then uploads and downloads may degrade call quality. AZTCO-FW software can prioritize the call traffic above other protocols, and ensure that the calls make it through clearly without breaking up, even while streaming hi-def video from Netflix at the same time. Instead of the call breaking up, the shaper reduces speed of the other transfers to leave room for the calls.

19.1.3 Reduce Gaming Lag

The shaper also has options to give priority to the traffic associated with network gaming. Similar to prioritizing VoIP calls, the effect is that even if users on the network are downloading while playing, the response time of the game should still be nearly as fast as if the rest of the connection were idle.

19.1.4 Keep P2P Applications In Check

By lowering the priority of traffic associated with known peer-to-peer ports, administrators can rest easier knowing that even if those programs are in use, they won't hinder other traffic on the network. Due to its lower priority, other protocols will be favored over P2P traffic, which will be limited when any other services need the bandwidth.

19.1.5 Enforce Bandwidth Limits

Limiters can apply a bandwidth limit to a group of devices, such as all traffic on an interface, or masking on limiters can apply them on a per-IP address or per-network basis. This way the firewall can ensure that no one person can consume all available bandwidth.

19.2 Hardware Limitations

Traffic shaping is performed with the help of ALTQ. Unfortunately, only a subset of all supported network cards are capable of using these features because the drivers must be altered to support ALTQ shaping. The following network cards are capable of using traffic shaping:

> ae(4), age(4), alc(4), ale(4), an(4), aue(4), axe(4), bce(4), bfe(4), bge(4), bridge(4), cas(4), cpsw(4), cxl(4), dc(4), de(4), ed(4), em(4), ep(4), epair(4), et(4), fxp(4), gem(4), hme(4), hn(4), igb(4), ix(4), jme(4), l2tp(4), le(4), lem(4), msk(4), mxge(4), my(4), ndis(4), nfe(4), ng(4), nge(4), npe(4), nve(4), ovpnc(4), ovpns(4), ppp(4), pppoe(4), pptp(4), re(4), rl(4), sf(4), sge(4), sis(4), sk(4), ste(4), stge(4), ti(4), tun(4), txp(4), udav(4), ural(4), vge(4), vlan(4), vmx(4), vr(4), vte(4), vtnet(4), xl(4)

Limiters use a different backend system, operating through dummynet pipes in ipfw and not through ALTQ. As such, all network cards may be used for Limiters, there are no restrictions. If a firewall contains a card that does not support ALTQ, it may use limiters instead.

19.3 Network Interface Drivers with ALTQ Traffic Shaping Support

The intention of this page is to provide information regarding FreeBSD's ALTQ drivers, what they do, and how they work.

19.3.1 Information

The ALTQ framework is used for queuing/traffic shaping. In AZTCO-FW® software, this is utilized by the Shaper Wizard and the Queues/Interfaces tabs under Firewall > Traffic Shaper.

See the altq(4) or the altq(9).

On that page, select the *version of FreeBSD that corresponds to the AZTCO-FW version being run*.

In addition to the drivers listed as supporting ALTQ in FreeBSD, AZTCO-FW software also includes support for ALTQ on vlan(4) and IPsec enc(4) interfaces.

If the NIC being used does not support ALTQ, *Limiters* may be used instead.

19.4 ALTQ Scheduler Types

AZTCO-FW® software contains several ALTQ scheduler types to cover a large range of shaping scenarios. The options for ALTQ are:

Priority Queuing (PRIQ) Manages prioritization of connections

Class-Based Queuing (CBQ) Supports bandwidth sharing between queues and bandwidth limits

Hierarchical Fair Service Curve (HFSC) Supports real-time bandwidth guarantees along with a hierarchical tree of nested queues.

Controlled Delay (CoDel) Attempts to combat bufferbloat.

Fair Queuing (FAIRQ) Attempts to fairly distribute bandwidth among all connections.

PRIQ, CBQ, and HFSC are selectable in the shaper wizards and the wizards will show the proper options and create the queues based on the chosen ALTQ discipline.

19.4.1 Performance Caveats

Enabling ALTQ traffic shaping places an extra burden on the hardware, and there will be an overall potential network performance loss. On systems that have horsepower to spare, this may not be noticeable. On systems that operate close to their specification limits the firewall may see a degradation of performance. Whether the loss is worse than working without shaping depends on the individual workload.

19.4.2 Priority Queuing (PRIQ)

PRIQ is one of the easiest disciplines to configure and understand. The queues are all directly under the root queue, there is no structure to have queues under other queues with PRIQ as there is with HFSC and CBQ. It does not care about bandwidth on interfaces, only the priority of the queues. The values for priority go from 0 to 15, and the higher the priority number, the more likely the queue is to have its packets processed.

PRIQ can be harsh to lesser queues, starving them when the higher priority queues need the bandwidth. In extreme cases, it is possible for a lower priority queue to have little or no packets handled if the higher priority queues are consuming all available resources.

19.4.3 Hierarchical Fair Service Curve (HFSC)

The HFSC traffic shaping discipline is very powerful. It is useful for services such as VoIP and video to deliver a minimum guaranteed amount of bandwidth.

Queues in HFSC are arranged in a hierarchy, or a tree, with root queues for each interface, parent queues underneath, and child queues nested under the parent queues (etc.). Each queue can have a set bandwidth and related options.

HFSC-specific Queue Options

HFSC supports a few queue options that are not supported by other disciplines. It is through these options that it achieves guaranteed real-time processing and link sharing.

The Service Curve (sc) is where bandwidth requirements for this queue are tuned. m1

> Burstable bandwidth limit d Time limit for bandwidth burst, specified in
>
> milliseconds. (e.g. 1000 = 1 second) m2 Normal bandwidth limit

For example, a connection needs m1 bandwidth within d time, but a normal maximum of m2. Within the initial time set by d, m2 is not checked, only m1. After d has expired, if the traffic is still above m2, it will be shaped. Most commonly, m1 and d are left blank, so that only m2 is checked.

Each of these values may be set for the following uses:

Upper Limit Maximum bandwidth allowed for the queue. Will do hard bandwidth limiting. The m1 parameter here can also be used to limit bursting. In the time frame d a connection will not get more than m1 bandwidth.

Real Time Minimum bandwidth guarantee for the queue. This is only valid for child queues. The m1 parameter will always be satisfied in time frame d, and m2 is the maximum that this discipline will allow to be used. Note The value for m2 cannot exceed 30% of the available bandwidth from the parent queue.

Link Share The bandwidth share of a backlogged queue. Will share bandwidth between classes if the Real Time guarantees have been satisfied. The m2 value for Link Share will override the Bandwidth setting for the queue. These two settings are the same, but if both are set, m2 from Link Share is used.

By combining these factors, a queue will get the bandwidth specified by the Real Time factors, plus those from Link Share, up to a maximum of Upper Limit. It can take a lot of trial and error, and perhaps a lot of arithmetic, but it may be worth it to ensure that network traffic is governed properly.

19.4.4 Class-Based Queuing (CBQ)

Class-Based Queuing, or CBQ, is similar to HFSC in that is can have a tree of queues nested under other queues. It supports bandwidth limits (not guarantees like HFSC), priorities for queues, and it has the ability to allow queues to borrow bandwidth from their parent. Because of the simpler queue configuration, it can be a good alternative to HFSC especially if the firewall does not need to guarantee minimum bandwidths.

With CBQ, queue priorities range from 0 to 7 with higher numbers indicating higher priority. Queues of an equal priority are processed in a round-robin fashion.

Note: Though child queues can borrow from their parent queue, the sum of the bandwidth of the child queues cannot exceed the bandwidth of the parent. Therefore, CBQ is not an alternative to limiters for individual (e.g. per-IP address) bandwidth limits.

CBQ-Specific Queue Options

The CBQ discipline supports the concept of *borrow*, meaning that if the Borrow from other queues when available checkbox on the queue is enabled, then the queue will be able to borrow other available bandwidth from its parent queue. This will only allow a child queue to obtain up to the bandwidth of its *immediate* parent, if available, it will not borrow from other parent queues.

19.4.5 CoDel Active Queue Management

The CoDel Active Queue Management (AQM) discipline is short for Controlled Delay and is pronounced "coddle". It was designed to combat problems associated with bufferbloat in networking infrastructure. Bufferbloat is described in detail at http://www.bufferbloat.net/projects/bloat/wiki/Introduction. Put simply, traffic can pile up and go in chunks rather than a smooth stream due to the size of buffers in network equipment. By controlling the delay of the traffic this effect can be lessened.

CoDel has no specific configuration controls or options. When activated for a queue, it will automatically attempt to manage traffic as described in the CoDel wiki at http://www.bufferbloat.net/projects/codel/wiki. It attempts to keep traffic delays low but does permit bursting, it controls delays but it does not pay attention to round-trip delay, load, or link speed, and it can automatically adjust if the link speed changes.

The target for CoDel is mid-range networking. It does not work well at very low bandwidth (1Mbit/s or less) and it does not gracefully handle large numbers of simultaneous flows or datacenter-grade traffic loads.

CoDel is not configurable using the wizard, but it does not require complex setup:

- Navigate to Firewall > Traffic Shaper, By Interface tab
- Select an interface (e.g. WAN)
- Set the Scheduler Type to *CODEL*
- Set an appropriate value for Bandwidth
- Click Save
- Repeat as needed for all other active WAN-type interface(s)

19.4.6 Fair Queuing (FAIRQ)

In FAIRQ, queues are monitored from highest to lowest priority, but the scheduler attempts to fairly distribute bandwidth among all connections.

When there is no contention for bandwidth, FAIRQ will send all waiting packets. When there is contention for bandwidth FAIRQ will start looking for queues that are not exceeding their limits, first starting with high priority queues and working toward lower queues. A packet in a full high priority queue is processed *after* a packet from a lower priority queue which is not full. If all queues are full, then FAIRQ will send a packet from the highest priority queue.

FAIRQ allows connections to exceed queue bandwidth, but will maintain an average consumption equal to the defined queue bandwidth.

FAIRQ is not currently supported in the traffic shaper wizard and it requires a manual configuration.

19.5 Advanced Customization

The rules and queues generated by the shaper wizard may not be an exact fit for a network. Network devices may use services that need shaped which are not listed in the wizard, games that use different ports, or other protocols that need limiting.

After the basic rules have been created by the wizard, it is relatively easy to edit or copy those rules to make adjustments for other protocols.

19.5.1 Editing Shaper Queues

Queues are where bandwidth and priorities are allocated by the shaper. Each queue has settings specific to the scheduler that was chosen in the wizard (*ALTQ Scheduler Types*). Queues can also be assigned other attributes that control how they behave. Queues may be managed at Firewall > Traffic Shaper. Click on a queue name in the list or tree shown on the By Interface or By Queue tabs, as seen in Figure *Traffic Shaper Queues List*

> Warning: Creating or editing queues is for advanced users only. It is a complex task with powerful results, but without thorough understanding of the settings involved the best practice is to stick with queues generated by the wizard rather than trying to make new queues.

To edit a queue, click its name in the list/tree.

To delete a queue, click it once to edit the queue, then click 🗑 Delete This Queue. Do not delete a queue if it is still being referenced by a firewall rule.

To add a new queue, click the interface or parent queue under which the new queue will be placed, and then click + Add New Queue.

When editing a queue, each of the options must be carefully considered. For more information about these settings than is mentioned here, visit the PF Packet Queuing and Prioritization FAQ or read *The OpenBSD PF Packet Filter* book.

Name The queue name must be between 1-15 characters and cannot contain spaces. The most common convention is to start the name of a queue with the letter "q" so that it may be more readily identified in the ruleset.

Priority The priority of the queue. Can be any number from 0-7 for CBQ and 0-15 for PRIQ. Though HFSC can support priorities, the current code does not honor them when performing shaping. Queues with higher numbers are preferred by the shaper when there is an overload, so situate queues accordingly. For example, VoIP traffic is the highest priority, so it would be set to a 7 on CBQ or 15 on PRIQ. Peer-to-peer network traffic, which can be delayed in favor of other protocols, would be set at 1.

Bandwidth (root queues) The amount of bandwidth available on this interface in the outbound direction. For example, WAN-type interface root queues list upload speed. LAN-type interfaces list the sum total of all WAN interface download bandwidth.

Queue Limit The number of packets that can be held in a queue waiting to be transmitted by the shaper. The default size is 50.

Scheduler Options There are five different Scheduler Options that may be set for a given queue:

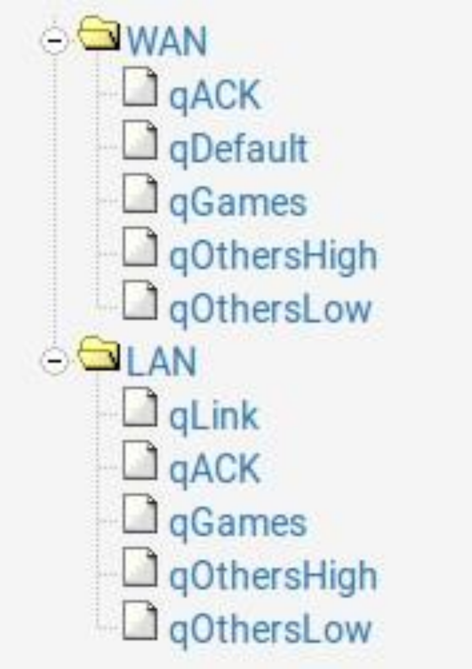

Fig. 1: Traffic Shaper Queues List

Default Queue Selects this queue as the default, the one which will handle all unmatched packets on an interface. Each interface must have one and only one default queue.

Random Early Detection (RED) A method to avoid congestion on a link. When set, the shaper will actively attempt to ensure that the queue does not get full. If the bandwidth is above the maximum given for the queue, drops will occur. Also, drops may occur if the average queue size approaches the maximum. Dropped packets are chosen at random, so connections using more bandwidth are more likely to see drops. The net effect is that the bandwidth is limited in a fair way, encouraging a balance. RED should only be used with TCP connections since TCP is capable of handling lost packets, and hosts can resend TCP packets when needed.

Random Early Detection In and Out (RIO) Enables RED with in/out, which results in having queue averages being maintained and checked against incoming and outgoing packets.

Explicit Congestion Notification (ECN) Along with RED, it allows sending of control messages that will throttle connections if both ends support ECN. Instead of dropping the packets as RED will normally do, it will set a flag in the packet indicating network congestion. If the other side sees and obeys the flag, the speed of the ongoing transfer will be reduced.

Codel Active Queue A flag to mark this queue as being the active queue for the Codel shaper discipline.

Description Optional text describing the purpose of the queue.

Bandwidth (Service Curve/Scheduler) The Bandwidth setting should be a fraction of the available bandwidth in the parent queue, but it must also be set with an awareness of the other neighboring queues. When using percentages, the total of all queues under a given parent cannot exceed 100%. When using absolute limits, the totals cannot exceed the bandwidth available in the parent queue.

Scheduler-specific Options Next are scheduler-specific options. They change depending on whether a queue is using HFSC, CBQ, or PRIQ. They are all described in *ALTQ Scheduler Types*.

Click Save to save the queue settings and return to the queue list, then click Apply Changes to reload the queues and activate the changes.

19.5.2 Editing Shaper Rules

Traffic shaping rules control how traffic is assigned into queues. If a new connection matches a traffic shaper rule, the firewall will assign packets for that connection into the queue specified by that rule.

Packet matching is handled by firewall rules, notably on the Floating tab. To edit the shaper rules:

- Navigate to Firewall > Rules
- Click the Floating Tab
- Find the rule to edit in the list, as shown in Figure *Traffic Shaper Rules List*
- Click ⧉ to edit an existing rule or ⊘ to create a copy of a rule
- Make any required adjustments to match different connections
- Save and Apply Changes as usual when editing firewall rules

Queues may be applied using *pass* rules on interface tabs, but the wizard only creates rules on the Floating tab using the *match* action that does not affect whether or not a connection is passed or blocked; it only queues traffic. Because these rules operate the same as any other rules, any criteria used to match connections may be used to queue.

See also:

For more information on floating rules, see *Floating Rules* and *Configuring firewall rules* for information on firewall rules in general.

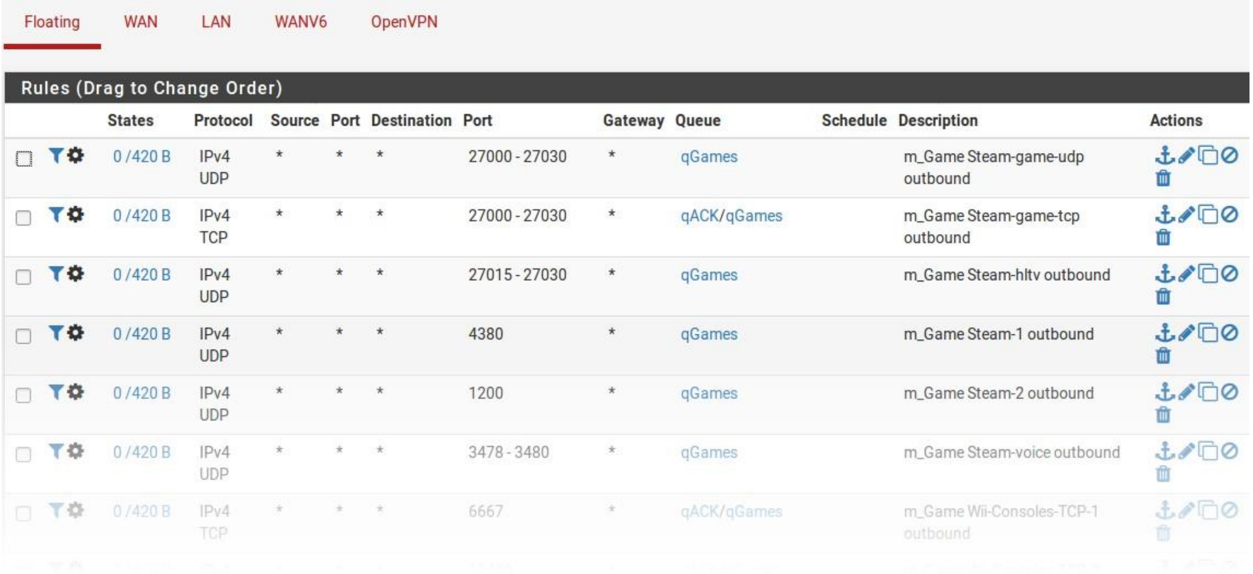

Rules (Drag to Change Order)

		States	Protocol	Source	Port	Destination	Port	Gateway	Queue	Schedule	Description	Actions
□		0 /420 B	IPv4 UDP	*	*	*	27000 - 27030	*	qGames		m_Game Steam-game-udp outbound	
□		0 /420 B	IPv4 TCP	*	*	*	27000 - 27030	*	qACK/qGames		m_Game Steam-game-tcp outbound	
□		0 /420 B	IPv4 UDP	*	*	*	27015 - 27030	*	qGames		m_Game Steam-hltv outbound	
□		0 /420 B	IPv4 UDP	*	*	*	4380	*	qGames		m_Game Steam-1 outbound	
□		0 /420 B	IPv4 UDP	*	*	*	1200	*	qGames		m_Game Steam-2 outbound	
□		0 /420 B	IPv4 UDP	*	*	*	3478 - 3480	*	qGames		m_Game Steam-voice outbound	
□		0 /420 B	IPv4 TCP	*	*	*	6667	*	qACK/qGames		m_Game Wii-Consoles-TCP-1 outbound	

Fig. 2: Traffic Shaper Rules List

Shaper Rule Matching Tips

Connections can be tricky to match properly due to several factors, including:

- NAT applies before outbound firewall rules can match connections, so for connections that have outbound NAT applies as they leave a WAN-type interface, the private IP address source is hidden by NAT and cannot be matched by a rule.
- Some protocols such as Bittorrent will use random ports or the same ports as other services.
- Multiple protocols using the same port cannot be distinguished by the firewall.
- A protocol may use a range of ports so wide that it cannot be distinguished from other traffic.

While many of these cannot be solved by the firewall directly, there are ways to work around these limitations in a few cases.

To match by a private address source outbound in WAN floating rules, first tag the traffic as it passes in on a local interface. For example, match inbound on LAN and use the advanced Tag field to set a value, and then use the Tagged field on the WAN-side floating rule to match the same connection as it exits the firewall. Alternately, queue the traffic as it enters the LAN with a pass rule instead of when it exits a WAN.

Match by address instead of port/protocol where possible to sort out ambiguous protocols. In these cases, either the local source or the remote destination may be a single address or a small set of addresses. For example, matching VoIP traffic is much simpler if the firewall can match the remote SIP trunk or PBX rather than attempting to match a wide range of ports for RTP (e.g. 10000- 20000).

If bittorrent is allowed on a network but must be shaped, then dedicate a specific local device that is allowed to use bittorrent and then shape all connections to/from that device as Peer-to-Peer traffic.

19.5.3 Removing Traffic Shaper Settings

To remove all traffic shaper queues and rules created by the wizard:

- Navigate to Firewall > Traffic Shaper
- Click the By Interface tab
- Click 🗑 Remove Shaper
- Click OK on the confirmation prompt

19.6 Traffic Shaping with Differentiated Services (DiffServ) Identifiers

AZTCO-FW® software supports Differentiated services (DiffServ) for traffic filtering or queue assignments. DiffServ takes the place of the outdated Type of service (TOS). DiffServ uses the upper six bits of the TOS field in the IP header (the six bits being called the *DiffServ Code Point field*), while the lower two bits are reserved for Explicit Congestion Notification (ECN).

Unless appropriately configured, AZTCO-FW software ignores the content of the DiffServ Code Point (DSCP) field. To prioritize traffic, the *Traffic Shaper* needs to be set up accordingly.

Warning: AZTCO-FW software *does not* support the setting or changing of DiffServ values, only matching.

19.6.1 Supported DiffServ Code Point Values

Note that the interpretations of the DSCP values, as provided by the various RFCs, are only given as a reference. How the DSCP values are interpreted in any specific setup is entirely up to the user or end nodes.

The Assured Forwarding (AF) Behavior Group is recommended in RFC 2597.

Table 1: Assured Forwarding (AF) Behavior Group values

Precedence	Class 1 (lowest)	Class 2	Class 3	Class 4 (highest)
Low Drop	AF11 (10/0x0a)	AF21 (18/0x12)	AF31 (26/0x1a)	AF41 (34/0x22)
Med Drop	AF12 (12/0x0c)	AF22 (20/0x14)	AF32 (28/0x1c)	AF42 (36/0x24)
High Drop	AF13 (14/0x0e)	AF23 (22/0x16)	AF33 (30/0x1e)	AF43 (38/0x26)

For low-drop/low-latency traffic, use EF and VA DSCP values.

Table 2: Expedited Forwarding (EF) and Voice Admit (VA) values

PHB	DSCP Value	RFC
Expedited Forwarding (EF)	46/0x2e	RFC 3246
Voice Admit (VA)	44/0x2c	RFC 5865

The Class Selector (CS) PHB group has been retained from TOS.

Table 3: Class Selector (CS) values

Class Selector	DSCP Value
CS1	8/0x08
CS2	16/0x10
CS3	24/0x18
CS4	32/0x20
CS5	40/0x28
CS6	48/0x30
CS7	56/0x38

To provide limited backward comparability to TOS, AZTCO-FW also recognizes the following DSCP/TOS values.

Table 4: TOS Compatibility values

TOS	DSCP Value	TOS value
reliability	1/0x01	4/0x04
throughput	2/0x02	8/0x08
lowdelay	4/0x04	16/0x10

AZTCO-FW only matches exact values. All six bit in the DSCP field must match.

19.6.2 Caveats

By default, AZTCO-FW matches only the first packet of a connection, which is the packet that creates an entry in the state table. If a connection starts with a different DSCP value, has no DSCP value in the starting packet, or otherwise changes DSCP values during the connection, the traffic will not be classified as expected.

Tip: This can be worked around by using “no state” rules, but crafting these rules in a secure manner is difficult, so it is not a workaround that we recommend.

19.6.3 Adding additional DSCP values for experimental use

Assuming basic knowledge about PHP, it is possible to add additional DiffServ Code Point values by editing /usr/local/www/guiconfig.inc. In this file, the variable $firewall_rules_dscp_types is initialized with an array containing the recognized DSCP values. New values can be specified as hex values, optionally followed by a blank and a comment like, for example:

```
"0x03",
```

Valid values are in the range 0x01 through 0x3f.

Caution: These changes will be lost upon a firmware update.

19.6.4 RFCs

- RFC 2474 — Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- RFC 2475 — An Architecture for Differentiated Services
- RFC 2597 — Assured Forwarding PHB Group
- RFC 2983 — Differentiated Services and Tunnels
- RFC 3086 — Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification
- RFC 3140 — Per Hop Behavior Identification Codes (replaces RFC 2836)
- RFC 3246 — An Expedited Forwarding PHB (Per-Hop Behavior) (obsoletes RFC 2598)
- RFC 3247 — Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior)
- RFC 3260 — New Terminology and Clarifications for Diffserv (updates
- RFC 2474, RFC 2475 and RFC 2597)
- RFC 4594 — Configuration Guidelines for DiffServ Service Classes
- RFC 5865 — A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic (updates RFC 4542 and RFC 4594)
- RFC 3289 — Management Information Base for the Differentiated Services Architecture
- RFC 3290 — An Informal Management Model for Diffserv Routers
- RFC 3317 — Differentiated Services Quality of Service Policy Information Base

19.7 Limiters

Limiters are an alternate method of traffic shaping. Limiters use dummynet(4) to enact bandwidth limits and perform other prioritization tasks, and they do not rely on ALTQ. Limiters are currently the only way to achieve per-IP address or per-network bandwidth rate limiting using AZTCO-FW® software. Limiters are also used internally by Captive Portal for per-user bandwidth limits.

Limiters are managed at Firewall > Traffic Shaper on the Limiters tab.

Like HFSC and CBQ, Limiters may be nested with queues inside other queues. Root-level limiters (Also called Pipes), may have bandwidth limits and delays, while child limiters (Also called queues), may have priorities (Also called weights). Bandwidth limits can be optionally masked by either the source or destination IP address, so that the limits can be applied on a per-IP address or network basis instead of as a general group.

Limiters are nearly always used in pairs: One for incoming traffic and one for outgoing traffic.

According to its man page the dummynet(4) system was originally designed as a means to test TCP congestion control and it grew up from there. Due to this purpose, a unique feature of limiters is that they can be used to induce artificial packet loss and delay into network traffic. That is primarily used in troubleshooting and testing (or being evil and playing a prank on someone), and not often found in production.

19.7.1 Uses for Limiters

The primary use for limiters is to apply bandwidth limits for users or specific protocols, e.g. “Maximum of 1Mbit/s for SMTP”, or “Joe’s PC only can use 5Mbit/s”. Limiters can apply a per-IP address or per-network limit, such as “All Users in 192.168.50.0/24 can use a maximum of 3Mbit/s each” or “The guest network and public network can use 1Mbit/s for each segment”.

Limiters are the only type of shaper available in AZTCO-FW software which is capable of oversubscription in this manner. The ALTQ shaper requires all child queues to sum up to no more than the speed of the parent queue, but masked limiters allow a set limit to as many IP addresses as can be funneled through the limiter by firewall rules.

Conceptually, consider a limiter as a bucket of bandwidth. All traffic flowing through an unmasked limiter draws bandwidth from the same bucket. Masking a limiter effectively sets up multiple buckets of the same size, one per masked group. Whether that is a single host or an entire network depends on the mask value.

Limiters can also allow for reserved bandwidth by limiting everything *except* a specific protocol which can then consume all remaining bandwidth. In this type of setup on a 10Mbit/s link the firewall would pass traffic from, for example, a SIP server with no limiter. Then the firewall would use a pass rule for all other traffic with a limit of 8Mbit/s. This would let the SIP server use all of the bandwidth it wanted, but it would always have a minimum of 2Mbit/s to itself.

19.7.2 How Limiters Work

Limiters, like ALTQ, hold traffic to a certain point by dropping or delaying packets to achieve a specific line rate. Usually taking advantage of built-in mechanisms from protocols that detect the loss and back off to a sustainable speed.

In situations where packets are queued under the same parent pipe, the firewall considers their weights when ordering the packets before it sends them. Unlike priorities in CBQ and PRIQ, the weight of a queue in a limiter will never starve it for bandwidth.

24.7.3 Limiters and IPv6

Limiters work with IPv6, though it requires separate IPv4 and IPv6 rules to apply limiters properly.

19.7.4 Limitations

Limiter pipes do not have a concept of borrowing bandwidth from other pipes. A limit is always a hard upper limit.

Limiters use IPFW, so there will be additional (though small) overhead from the IPFW kernel module and the extra packet processing involved.

Limiters cannot effectively guarantee a minimum bandwidth amount for a pipe or queue, only a maximum.

Child queues cannot have bandwidth values, so a pipe cannot be split into smaller pipes by queues. Child queues can only use weights to prioritize packets inside a pipe.

The overhead from delaying and queuing packets can cause increased mbuf usage. For more information on increasing the amount of available mbufs, see *Hardware Tuning and Troubleshooting*.

Limiters and Multi-WAN

When using limiters with Multi-WAN, limits for non-default gateways must be applied using floating rules set for the *out* direction and configured with the appropriate gateway.

19.7.5 Creating Limiters

Limiters are managed under Firewall > Traffic Shaper on the Limiters tab.

To create a new root-level limiter (pipe), click + New Limiter.

To create a child limiter (queue), click an existing limiter under which it can be created, and click + Add New Queue.

Tip: In nearly all cases, limiters exist in pairs at the same level (e.g. two pipes, or two queues): One for inbound traffic and one for outbound traffic. When creating new limiters or queues, create one for each direction.

Enable Check the box to enable this limiter. If the limiter is disabled, it will not be available for use by firewall rules.

Name This defines the name of the limiter, as it will appear for selection on firewall rules. The name must be alphanumeric, and may also include - and _.

> Tip: When choosing a name, avoid using In and Out since the same limiter, if used on both WAN and LAN, would be used in the *In* direction on one interface and the *Out* direction on another. The best practice is to use Down or Download and Up or Upload.

Bandwidth (Pipes) This section defines a bandwidth value for the pipe, or multiple bandwidths if schedules are involved. This option does not appear when editing a child limiter (queue).

Bandwidth The numerical part of the bandwidth for the pipe, e.g. 3 or 500.

24.7. Limiters

Bw Type The units for the Bandwidth field, such as *Mbit/s*, *Kbit/s*, or *Bit/s*.

Schedule If the firewall has schedules defined (*Time Based Rules*), the firewall offers them in this list. When schedules are in use by the firewall, the limiter can have a bandwidth value for each potential schedule. Define these by clicking + Add Schedule to add another bandwidth definition.

If a limiter contains multiple bandwidth specifications, they must each use a different schedule. For example if the firewall has a "Work Day" schedule, then it must also have an "Off Hours" schedule that contains all of the time not included in "Work Day" for the second bandwidth specification.

Mask This drop-down list controls how the limiter will mask addresses in the pipe or queue.

None When set to *none*, the limiter does not perform any masking. The pipe bandwidth will be applied to all traffic as a whole.

Source / Destination address When a limiter is set for *Source Address* or *Destination Address*, the pipe bandwidth limit will be applied on a per-IP address basis or a subnet basis, depending on the masking bits, using the direction chosen in the masking.

In general, a limiter should mask the Source Address on Upload (In) limiters for LAN-type interfaces, and Destination Address on Download (Out) limiters on LANtype interfaces. Similar to swapping the directionality of the limiters when applying to LAN and WAN, masking is swapped as well, so the same masked limiter set for In on LAN should be used for Out on WAN.

Mask Bits There are separate boxes to control the address masking for IPv4 and IPv6. For IPv4 a value of *32* for IPv4 mask bits sets up a per-IPv4 address limit, which is the most common usage. For a per-IPv6-address limit, use *128* as the IPv6 mask bits value.

To create per-subnet or similar masks, enter the subnet bits in the appropriate field for either IPv4 or IPv6 mask bits, such as *24* to limit IPv4 in groups of /24 subnets.

Description An optional bit of text to explain the purpose for this Limiter.

Advanced Options Additional options that vary when editing a pipe or a queue.

Delay (Pipes) The Delay option is only found on limiter pipes. It introduces an artificial delay (latency), specified in milliseconds, into the transmission of any packets in the limiter pipe. This is typically left blank so that packets are transmitted as fast as possible by the firewall. This can be used to simulate high-latency connections such as satellite uplinks for lab testing.

Weight (Queues) The Weight option is only found on child limiters (queues). This value can range from 1 to 100. Higher values give more precedence to packets in a given queue. Unlike PRIQ and CBQ priorities, a lowly-weighted queue is not in danger of being starved of bandwidth by the firewall.

Packet loss rate Another method of artificially degrading traffic. The Packet Loss Rate can be configured to drop a certain fraction of packets that enter the limiter. The value is expressed as a decimal representation of a percentage, so 0.01 is 1%, or one packet out of a hundred dropped. This field is typically left empty so every packet is delivered by the firewall.

Queue Size Sets the size of the queue, specified in queue slots, used for handling queuing delay. Left blank, it defaults to 50 slots, which is the recommended value. Slow speed links may need a lower queue size to operate efficiently. High speed links may need more slots.

24.7. Limiters

Tip: In cases where there are several limiters or limiters with large Queue Size values, a System Tunable may need set to increase the value of net.inet.ip. dummynet.pipe_slot_limit above the total number of configured queue lots among all pipes and queues.

Bucket Size The Bucket Size, also specified in slots, sets the size of the hash table used for queue storage. The default value is 64. It must be a numeric value between 16 and 65536, inclusive. This value is typically left blank.

See also:

For more information about these values, consult the ipfw(8) man page, in the section titled “Traffic Shaper (Dummynet) Configuration”.

19.7.6 Assigning and Using Limiters

Limiters are assigned using firewall rules via the In/Out Pipe selectors under Advanced Options. Any potential matching criteria that a firewall rule supports can assign traffic to a limiter.

The most important thing to remember when assigning a limiter to a rule is that the In and Out fields are designated from the perspective of the firewall itself.

For example, in a firewall configuration with a single LAN and single WAN, inbound traffic on a LAN interface is leaving toward the Internet, i.e. *uploaded* data. Outbound traffic on the LAN interface is going toward the client PC, i.e. *downloaded* data. On the WAN interface the directionality is reversed; Inbound traffic is coming from the Internet to the client (download), and outbound traffic is going from the client to the Internet (upload).

In most cases, a firewall rule will have both an In limiter and Out limiter, but only the In limiter is required by the firewall to limit traffic in a single direction.

Limiters may be applied on normal interface rules, or on floating rules. On floating in the *out* direction, the In/Out selections are flipped conceptually.

19.7.7 Checking Limiter Usage

Information about active limiters may be found under Diagnostics > Limiter Info. Here, each limiter and child queue is shown in text format.

The set bandwidth and parameters for each limiter are displayed by the page, along with the current traffic level moving inside the limiter. In the case of masked limiters, the firewall displays the bandwidth of each IP address or masked group.

19.8 Traffic Shaping and VPNs

The following discussions pertain primarily to ALTQ shaping. Limiters will work fine with VPNs as they would with any other interface and rules. Only the ALTQ shaper requires special consideration.

Traffic shaping with VPNs is a tricky topic because VPN traffic is considered separate from, but also a part of, the WAN traffic through which it also flows. If WAN is 10 Mbit/s, then the VPN can also use 10Mbit/s, but there is not actually 20Mbit/s of bandwidth to consider, only 10Mbit/s. As such, methods of shaping that focus more on prioritization than bandwidth are more reliable, such as PRIQ or in some cases, CBQ.

If all traffic inside the VPN must be prioritized by the firewall, then it is enough to consider only the VPN traffic itself directly on WAN, rather than attempting to queue traffic on the VPN separately. In these cases, use a floating rule on

24.8. Traffic Shaping and VPNs
WAN to match the VPN traffic itself. The exact type of traffic varies depending on the type of VPN. IPsec and PPTP traffic on WAN can both be prioritized by the shaper wizard, and these rules can be used as an example to match other protocols.

19.8.1 OpenVPN

With OpenVPN, multiple interfaces exist on the operating system, one per VPN. This can make shaping easier in some cases. Features of OpenVPN can also make it easier to shape traffic on WAN and ignore the tunnel itself.

Shaping inside the tunnel

If multiple classes of traffic are carried on the tunnel, then prioritization must be done to the traffic inside the tunnel. In order for the wizard to consider the traffic in this way, the VPN must be assigned as its own interface in the GUI. To accomplish this, assign it as described in *Interface assignment and configuration*, and then use the shaper wizard as if it were a separate WAN interface, and classify the traffic as usual.

Shaping outside the tunnel (passtos)

If the primary concern is shaping VoIP traffic over a VPN, another choice to consider is the passtos option in OpenVPN, called Type-of-Service in the OpenVPN client or server options. This option copies the TOS bit from the inner packet to the outer packet of the VPN. Thus, if the VoIP traffic has the TOS (DSCP) portion of the packet header set, then the OpenVPN packets will also have the same value.

This option is more useful for signaling intermediate routers about the QoS needs, however. Though the DSCP option on firewall rules can match based on TOS bits, as described in *Diffserv Code Point*, such matching would have to occur in the packet creating a firewall state, and not on specific packets flowing through that state.

Note: Because this option tells OpenVPN to copy data from the inner packet to the outer packet, it does expose a little information about the type of traffic crossing the VPN. Whether or not the information disclosure, though minor, is worth the risk for the gains offered by proper packet prioritization depends on the needs of the network environment.

19.8.2 IPsec

IPsec is presented to the operating system on a single interface no matter how many tunnels are configured and no matter which WANs are used by the tunnels. This makes shaping IPsec traffic difficult, especially when trying to shape traffic inside one particular IPsec tunnel.

The IPsec interface is also not possible to use on its own as an interface with the wizard. Floating rules can match and queue traffic on the IPsec interface, but in most cases only inbound traffic will be queued as expected. Actual results may vary.

24.8. Traffic Shaping and VPNs

19.9 Traffic Shaping UPnP Connections

UPnP rules are generated dynamically by the UPnP daemon happen outside of the typical user rules. There is a configuration option for UPnP where a queue can be defined to which UPnP will direct traffic that is directed through the rules it creates. This may be set through the AZTCO-FW® webGUI at Services > UPnP & NAT-PMP, and type in a valid Traffic Shaping Queue.

See also:

Monitoring the Queues Using the Shaper Wizard to Configure ALTQ Traffic Shaping Troubleshooting Traffic Shaping Troubleshooting Traffic Shaping Graphs

Traffic shaping, or network Quality of Service (QoS), is a means of prioritizing network traffic. Without traffic shaping, packets are processed on a first in/first out basis by the firewall. QoS offers a means of prioritizing different types of traffic, ensuring that high priority services receive the bandwidth they need before lesser priority services.

For simplicity, the traffic shaping system in AZTCO-FW® software may also be referred to as the “shaper”, and the act of traffic shaping may be called “shaping”.

19.10 Traffic Shaping Types

There are two types of QoS available in AZTCO-FW software: ALTQ and Limiters.

The ALTQ framework is handled through pf and is closely tied to network card drivers. ALTQ can handle several types of schedulers and queue layouts. The traffic shaper wizard configures ALTQ and gives firewall administrators the ability to quickly configure QoS for common scenarios, and it allows custom rules for more complex tasks. ALTQ is inefficient, however, so the maximum potential throughput of a firewall is lowered significantly when it is active.

AZTCO-FW software also supports a separate shaper concept called Limiters. Limiters enforce hard bandwidth limits for a group or on a per-IP address or network basis. Inside of those bandwidth limits, limiters can also manage traffic priorities.

19.11 Traffic Shaping Basics

For administrators who are unfamiliar with traffic shaping, it is like a bouncer at an exclusive club. The VIPs (Very Important Packets) always make it in first and without waiting. The regular packets have to wait their turn in line, and “undesirable” packets can be kept out until after the real party is over. All the while, the club is kept at capacity and never overloaded. If more VIPs come along later, regular packets may need to be tossed out to keep the place from getting too crowded.

ALTQ shaping concepts can be counter-intuitive at first because the traffic has to be queued in a place where the operating system can control the flow of packets. Incoming traffic from the Internet going to a host on the LAN (downloading) is shaped *leaving* the LAN interface from the firewall. In the same manner, traffic going from the LAN to the Internet (uploading) is shaped when leaving the WAN.

For ALTQ, there are traffic shaping queues, and traffic shaping rules. The queues allocate bandwidth and priorities. Traffic shaping rules control how traffic is assigned into those queues. Rules for the shaper work the same as firewall rules, and allow the same matching characteristics. If a packet matches a shaper rule, it will be assigned into the queues specified by that rule. In AZTCO-FW software, shaper rules are mostly handled on the Floating tab using the *Match* action that assigns the traffic into queues, but rules on any interface can assign traffic into queues using the *Pass* action.

Limiter rules are handled differently. Limiters apply on regular pass rules and enforce their limits on the traffic as it enters and leaves an interface. Limiters almost always exist in pairs: One for the “download” direction traffic and one for the “upload” direction traffic.

CHAPTER

TWENTY

CAPTIVE PORTAL

20.1 Captive Portal Zones

Captive Portal zones define separate portals for different sets of interfaces. For example, LAN and Wireless could use one portal, while a conference room would get a separate portal page. Each zone has separate settings for HTML pages, authentication, allowed addresses, and so on. A zone must be created before its settings can be changed.

Note: A zone may have multiple interfaces, but an interface may only be a member of one zone. Attempting to add the same interface to multiple zones will result in an error.

20.1.1 Managing Captive Portal Zones

Captive Portal zones are managed at Services > Captive Portal. A list of zones is displayed there, and zones may be added, edited, or deleted from that list.

To create a new Captive Portal zone:

- Navigate to Services > Captive Portal
- Click ＋ Add
- Enter a Zone Name, which may only consist of letters, digits, numbers, and underscores. Spaces and other special characters may not be used
- Enter an optional Zone Description to further describe the zone, if desired
- Click Save & Continue to move on to the portal settings for the zone To edit an existing zone, click ⧉ at the end of its row.

To delete an existing zone, 🗑 clickat the end of its row, and then click to confirm the action.

20.2 Common Captive Portal Scenarios

The following are some basic, common scenarios for the use of a Captive Portal. The details of how to perform all of the actions described will be covered throughout this chapter.

20.2.1 Portal Configuration Without Authentication

For a simple portal without authentication:

- Create a new Zone
- Check Enable captive portal
- Select an Interface
- Upload an HTML page with the portal contents as described in *Portal page without authentication*
- Click Save

Additional configuration options may be added as detailed in *Zone Configuration Options*.

20.2.2 Portal Configuration Using Local Authentication or Vouchers

To setup a portal with local authentication:

- Create a Zone
- Check Enable captive portal
- Select an Interface
- Set Authentication Method to Local User Manager / Vouchers
- Upload an HTML page with the portal contents as described in *Portal page with authentication*.

Additional configuration options may be added as detailed in *Zone Configuration Options*. Then configure the local users in the User Manager (*User Management and Authentication*).

To use vouchers, proceed to the Vouchers tab and create them there. See *Vouchers* for more information on Vouchers, and use the sample portal page HTML code from *Portal page with Vouchers*.

20.2.3 Portal Configuration Using RADIUS Authentication

To setup a portal using RADIUS authentication:

- Configure the RADIUS server to allow requests from the firewall
- Create a Zone
- Check Enable captive portal
- Select an Interface
- Set Authentication Method to RADIUS Authentication
- Fill in the settings for Primary RADIUS Server under Primary Authentication Source

Read the next section for information on specific configuration options.

20.3 Zone Configuration Options

This section describes each of the configuration options for a Captive Portal zone. Options for a zone are independent of those for other zones. For example, allowed IP address entries in a zone only affect that specific zone.

To reach this page, navigate to Services > Captive Portal and edit an existing zone from the list with ⧉, or click + Add to create a new zone.

Enable Check to enable this Captive Portal zone.

Description Brief text describing the purpose of the zone.

Interface Determines the interfaces that used by this Captive Portal zone. This cannot be a WAN interface. It can be a bridge interface so long as it is the actual bridge (e.g. bridge0) and the bridge interface has an IP address assigned.

Maximum concurrent connections Specifies the maximum number of concurrent connections to the portal web server per IP address. The default value is 4, which is sufficient for most environments. This limit exists to prevent a single host from exhausting all resources on the firewall, whether inadvertent or intentional.

One example where this would otherwise be a problem is a host infected with a worm. The thousands of connections issued will cause the captive portal page to be generated repeatedly if the host is not authenticated already, which would otherwise generate so much load it could leave the firewall unresponsive.

Idle timeout A timeout, specified in minutes, after which idle users will be disconnected by the portal. Users may log back in immediately.

Hard timeout A timeout, specified in minutes, after which the portal will forcefully log off users.

Tip: Set either a hard timeout, idle timeout, or both to ensure sessions are removed by the portal when users do not log off manually.

Users may log back in immediately after the hard timeout if their credentials are still valid (for local accounts, not expired, and for RADIUS authentication, user can still successfully authenticate to RADIUS).

Note: If a timeout value is set, the timeout must be less than the DHCP lease time or captive portal sessions can remain active for IP addresses that have switched to different devices. Setting the timeout lower will ensure that the portal sessions end before the lease would be reallocated to a new client.

Traffic Quota An amount of traffic which, when exceeded by a client, will trigger a disconnect of that client by the portal. This includes both upload and download traffic. Users may log back in immediately if their credentials are still valid

Pass-Through Credits These credits give devices a grace period before they must authenticate via the portal. For example, a device could connect 3 times within a day without seeing the portal page, but any more than that and they must login. By setting the hard timeout to a value such as 1 hour, the portal would effectively limit a client to three hours of access before forcing it to authenticate. By default this is disabled, and all clients are presented with the portal login page and must login.

Note: For this to be effective, set a hard timeout and/or idle timeout.

Pass-through credits allowed per MAC address The number of times a specific MAC address may connect through the portal. Once the client uses its credits, it can only log in with valid credentials until the waiting period has expired.

Waiting period to restore pass-through credits The number of hours after which the portal will restore the pass-through credits for a client to the original count after it uses the first one. This must be above 0 hours.

Reset waiting period on attempted access If enabled, the waiting period is reset by the portal to the original duration if access is attempted when all pass-through credits have already been exhausted. This prevents people who repeatedly attempt to access the portal from gaining open access too quickly.

Logout popup window When checked, the portal attempts to show a logout pop up window to the user which allows clients to explicitly disconnect themselves before the idle or hard timeout occurs. Unfortunately, since most browsers block pop up windows, this window may not work for most clients unless.

Pre-authentication redirect URL As the name implies, this option redirects users to the specified URL *before* they authenticate. Commonly, this is used to display a custom landing page describing the device location hosted on a server locally or elsewhere. That landing page must contain a link which in turn redirects the users back to the portal page, e.g. http://x.x.x.x:8002/index.php?zone=somezone&redirurl=http%3A%2F%2Fsomesite.example.com.

See also:

See *Allowed Hostnames* to allow hostnames through the portal without authentication, and *Allowed IP Address* for IP addresses.

The custom captive portal page must have extra code at the top to properly handle this redirect. In the example code below, the pre-authentication redirect target page must also put its own URL in the redirurl parameter of its link back to the portal in order for the login page to appear.

```
<?php require_once("globals.inc");
$request_uri = urldecode(str_replace("/index.php?zone={$_REQUEST['zone']}&redirurl=", ,→"", $_SERVER["REQUEST_URI"]));
$portal_redirurl = urldecode("$PORTAL_REDIRURL$");
if(!stristr($portal_redirurl, $request_uri)) { Header("Location:
$PORTAL_REDIRURL$"); exit;
}
?>
```

After authentication Redirection URL After authenticating or clicking through the portal, it will redirect users to this URL rather than the one they originally tried to access. If this field is left blank, the portal will redirect the user to the URL they initially attempted to access.

Blocked MAC address redirect URL URL to which the portal will redirect users with blocked MAC addresses when they attempt access through the portal.

Preserve user database When set, the database containing logged-in users is preserved by the portal when the firewall reboots.

Concurrent user logins Controls whether or not users are allowed to connect multiple times. This is not a total limit for the entire portal, but a per-account limit.

May be set to one of the following:

Disabled The portal will not allow concurrent logins for a user or voucher.

Multiple (Default) The portal does not enforce any restrictions on concurrent logins by a user or voucher.

Last Login The portal will only allow only one login per user account or voucher. The most recent login is permitted and any previous logins are disconnected.

First Login The portal will only allow only one login per user account or voucher. The portal permits the first login and denies any subsequent login attempt.

MAC filtering When set, the portal disables MAC address filtering. This is necessary in cases where the MAC address cannot reliably be determined, such as when multiple subnets exist behind a separate router using the portal. In that type of situation, all users behind a router will show up to the portal with the MAC address of the intermediate router. If this option is set, the portal will not attempt to ensure that the MAC address of clients stay the same while they are logged into the portal.

Note: This option is not compatible with RADIUS MAC authentication.

Pass-through MAC Auto Entry In certain use cases, users may only need to authenticate once per device, and then the client should not see the portal login again unless they change devices. Setting up pass-through MAC entries can automatically achieve this goal.

Pass-through MAC automatic additions If this option is set, the portal automatically adds a MAC passthrough entry after the user has successfully authenticated. Users of that MAC address will never have to authenticate again unless so long as the entry is present in the configuration. To remove the passthrough MAC entry, log in and remove it manually from the Pass- through MAC tab.

Note: This option is not compatible with RADIUS MAC authentication or the logout window.

Pass-through MAC automatic addition with username If this option is set, the portal saves the username used during authentication along with the pass-through MAC entry. To remove the passthrough MAC entry, log in and remove it manually from the Passthrough MAC tab.

Per-user bandwidth restrictions Captive Portal can also optionally rate-limit users to keep them from using too much bandwidth. The Default download and Default upload fields define the default values for user bandwidth, specified in Kilobits per second. These values can be overridden by RADIUS (*Passing back configuration from RADIUS Servers*) for different limits for specific users. If the fields are blank or set to 0, then users have unlimited bandwidth.

Use Custom Captive Portal Page When set, enables upload controls for manually crafted portal pages. See *HTML Page Contents* for details. When unset, simple portal page customization controls are available (*Captive Portal Login Page*).

20.3.1 Captive Portal Login Page

These simple customization controls enable small changes to the portal page without writing custom HTML. For more complicated portal pages, see *HTML Page Contents*.

Display Custom Logo Image When set, the portal page includes the custom image from Logo Image instead of the default logo.

Logo Image Upload control for setting a custom logo image.

Display Custom Background Image When set, the portal page includes the custom image from Background Image instead of the default background.

Background Image Upload control for setting a custom background image.

Terms and Conditions Text displayed by the portal to the user to which the user must agree before they are permitted to login.

20.3.2 Authentication

This section configures authentication for Captive Portal. If authentication is required for the zone it may be handled by the local user database, RADIUS, or LDAP.

Authentication Method

Use an Authentication backend This option allows users to authenticate with a username and password or vouchers. The authentication is handled by the local user database (*User Management and Authentication*) or an authentication server (*Authentication Servers*).

Vouchers are pre-generated access codes which grant short-term access to users. Vouchers may be used in addition to, or instead of, user authentication. For more information on using vouchers, see *Vouchers* later in this section.

None, don’t authenticate users The portal only requires users to click through the login page for access. The form must still be submitted, but it does not need to have any user entry fields, only a submit button.

Use RADIUS MAC Authentication The portal attempts to authenticate users by sending their MAC address as the username and the password entered into MAC authentication secret to the RADIUS server.

Note: Users must still attempt an HTTP connection so the portal will see the attempt and perform the initial authentication.

See *RADIUS MAC Authentication Options* for additional options.

This option is not available if MAC filtering is disabled.

Authentication Server A multi-select control where one or more primary authentication servers, or the local database, can be set for use by the portal. See *Primary Authentication Source* for more information.

Local Database Captive Portal users in this mode are managed in the AZTCO-FW® GUI. Local users are added in the User Manager (*Manage Local Users*).

Additionally, the Local Authentication Privileges option can limit access to only users who possess the proper access privileges.

LDAP Server When an LDAP server is active in the control, it is used by the portal for authentication as-is. There are no additional options for LDAP server behavior.

RADIUS Server When a RADIUS server is active in the control, numerous RADIUS server options are displayed by the GUI and Captive Portal users in this zone will be validated against the configured RADIUS server(s).

Secondary Authentication Server Similar to Authentication Server, but sets up an additional separate means of authentication using distinct fields. See *Secondary Authentication Source* for more information.

Primary Authentication Source

The Primary/Secondary authentication servers are used for the main username and password fields on the login form, auth_user and auth_pass, such as:

```
<tr>
        <td align="right">Username:</td>
        <td><input name="auth_user" type="text" style="border: 1px dashed;"></td>
  </tr> <tr>
        <td align="right">Password:</td>
        <td><input name="auth_pass" type="password" style="border: 1px dashed;"></td> </tr>
```

If the first server is down, the portal will attempt authentication using the other servers in the list, in order (top down).

Secondary Authentication Source

The secondary authentication source defines a completely separate authentication setup from the primary. For example, the primary source could be traditional usernames and passwords, while the secondary could be pre-paid card numbers or PINs.

The secondary authentication source uses the form fields auth_user2 and auth_pass2 in the captive portal HTML, such as:

```
<tr>
        <td align="right">Username:</td>
        <td><input name="auth_user2" type="text" style="border: 1px dashed;"></td>
</tr> <tr>
        <td align="right">Password:</td>
        <td><input name="auth_pass2" type="password" style="border: 1px dashed;"></td> </tr>
```

If the first server is down, the portal will attempt authentication using the other servers in the list, in order (top down).

Local Database Options

Local Authentication Privileges When Allow only users/groups with 'Captive portal login' privilege set is active, the portal will limit access to only users who have Captive Portal privilege. The privilege can be directly on their account or inherited via group membership..

RADIUS Authentication Options

RADIUS is a means of authenticating users against a central server that contains account information. There are many implementations of RADIUS, such as FreeRADIUS, Radiator, and NPS on Windows servers.

RADIUS accounting can be enabled to send usage information for each user to the RADIUS server. Refer to documentation for the RADIUS server for more information.

See also:

To add or edit RADIUS server entries on AZTCO-FW software, see *Authentication Servers*.

See also:

For those with a Microsoft Active Directory network infrastructure, RADIUS can be used to authenticate captive portal users from Active Directory using Microsoft NPS. This is described in *Authenticating from Active Directory using RADIUS/NPS*.

Passing back configuration from RADIUS Servers

Some default Captive Portal settings can be overridden by reply attributes from RADIUS servers. The exact attributes can vary by vendor, and may not be supported by all RADIUS servers.

User bandwidth restrictions Defines the bandwidth for the user, drawn from common options such as WISPr-Bandwidth-Max-Up/WISPr-Bandwidth-Max- Down, or ChilliSpot-Bandwidth-Max-Up/ChilliSpot-Bandwidth-Max-Down.

Session Timeout Drawn from the RADIUS attribute Session-Timeout, it will disconnect the user after the time specified by the RADIUS server.

Idle Timeout Drawn from the RADIUS attribute Idle-Timeout, it will disconnect the user after the time specified by the RADIUS server.

Accounting Interval Interim Taken from Acct-Interim-Interval, it directs the portal to send interim accounting updates at the specified interval.

URL Redirection Allows the after-authentication redirect URL to be defined by the RADIUS server through WISPr-Redirection-URL.

RADIUS Options

These options fine-tune how RADIUS authentication behaves.

NAS Identifier Configures an alternate NAS Identifier to send with RADIUS requests. The default value is CaptivePortal-<zone name>.

Reauthentication If enabled, the portal sends Access-Request packets to the RADIUS server for each user that is logged in every minute. If an Access-Reject is received for a user, that user is disconnected from the captive portal immediately. This allows actively terminating user sessions from the RADIUS server.

> Warning: If concurrent login limits are defined in RADIUS this option may not work properly, as the additional request would fail as the reauthentication attempt would be considered a second concurrent login.

> Note: If reauthentication is combined with RADIUS accounting, Interim accounting updates must be used to track usage during sessions, otherwise the RADIUS server will not know if a user exceeds limits until they logout.

Session-Timeout When set, clients will be disconnected after the amount of time set by the RADIUS Session-Timeout attribute sent to the portal at login.

Traffic Quota When set, the portal uses the AZTCO-FW-Max-Total-Octets reply attribute sent by the RADIUS server to set a traffic quota for a user. This determines an amount of traffic which, when exceeded by a client, will trigger a disconnect of that client by the portal. This includes both upload and download traffic.

Per-user Bandwidth Restrictions When set, the portal uses the AZTCO-FW-Bandwidth-Max-Up and AZTCO-FW-Bandwidth-Max-Down reply attribute sent by the RADIUS server to set per-user bandwidth restrictions.

MAC address format This option changes the MAC address format used in RADIUS. Change this to alter the username format for RADIUS MAC authentication to one of the following styles: Default Colon-separated pairs of digits: 00:11:22:33:44:55

Single Dash Digits in two groups, separated by a single dash halfway: 001122-334455

IETF Hyphen-separated pairs of digits: 00-11-22-33-44-55

Cisco Groups of four digits separated by a period: 0011.2233.4455

Unformatted All digits together with no formatting or separators: 001122334455

RADIUS MAC Authentication Options

RADIUS MAC Secret When the portal attempts RADIUS MAC authentication, it sends the MAC Address as the username and this value as the password.

Login Page Fallback When set, the portal will redirect a client to the login page if MAC Authentication failed.

Accounting

RADIUS accounting sends session information back to the RADIUS server indicating when a user session starts, ends, and how much data they have transmitted.

Warning: Not all RADIUS servers support or are configured to accept accounting data. Setup the RADIUS server properly before enabling this feature.

Accounting Server An authentication server entry for a RADIUS server to which the portal will send accounting data (*Authentication Servers*).

Send Accounting Updates Configures the specific type of accounting supported by the server.

No updates Synonymous with disabling accounting, the portal will not send accounting updates to the server.

Stop/start The portal sends START and STOP records for a user session only.

Stop/start (FreeRADIUS) The portal sends START and STOP records for a user session only, in a way that is compatible with FreeRADIUS.

Interim The portal sends START and STOP records and also periodically sends updates to the server while a user session is active. This method is less likely to lose session data if the firewall restarts without notifying the RADIUS server of a STOP message, but will cause increased database usage on the RADIUS server.

Accounting Style When Invert Acct-Input-Octets and Acct-Output-Octets is enabled, data counts for RADIUS accounting packets will be taken from the client perspective, not the NAS.

Acct-Input-Octets will represent download, and Acct-Output-Octets will represent upload.

Idle Time Accounting This option changes the time sent in the STOP message for a user disconnected by the portal for idle timeout. When unset (default), the time sent is the last activity time. When set, the idle time is included.

20.3.3 HTTPS login

Login When set, the portal will listen for and accept HTTPS requests for the portal page. This option requires an SSL/TLS Certificate.

HTTPS server name The FQDN (hostname + domain) used by the portal for HTTPS. This must match the Common Name (CN) on the certificate to prevent users from receiving certificate errors.

SSL Certificate Select the SSL certificate used by the portal for HTTPS . Certificates are managed in *Certificate Management*.

Disable HTTPS Forwards When checked, attempts by clients to connect to HTTPS sites on port 443 are not redirected to the portal. This prevents users from receiving invalid certificate errors. Users must attempt a connection to an HTTP site, and will then be forwarded to the portal.

20.3.4 HTML Page Contents

When Use custom captive portal page is set on the zone, the portal displays these controls to upload custom HTML pages to alter the look of the page presented to users when they are redirected to the portal.

Customizing these pages is optional. Any page contents left blank will use internal defaults.

Portal pages may contain PHP code, and may also include other resources such as images and CSS files. See *File Manager* for more information on including additional assets in a custom portal page.

> Warning: Since custom portal pages can run PHP, audit the code to ensure security so the page cannot be exploited by connecting users. Also, avoid granting privileges to this page to untrusted administrators.

In each individual section, the pages can be managed by the displayed controls:

- To upload a new page, click Browse and select the file to upload. When the portal options are saved, the file will be copied.
- To view an existing page, click 🗎 View Page Contents
- To download a copy of an existing page, click ⬇ Download
- To erase the custom page, click ↺ Restore Default Page

Portal page contents

This control is for the main portal page presented to users. Depending on the selected options for the portal, use one of the following examples as the basis for a custom page.

Portal page without authentication

This shows the HTML of a portal page that can be used without authentication.

```
<html>
<head>
        <title>Welcome to our portal</title>
</head> <body>
        <p>Welcome to our portal</p>
        <p>Click Continue to access the Internet</p>
        <form method="post" action="$PORTAL_ACTION$">
                <input name="redirurl" type="hidden" value="$PORTAL_REDIRURL$">
                      <input name="zone" type="hidden" value="$PORTAL_ZONE$">
                <input name="accept" type="submit" value="Continue">
        </form>
</body>
</html>
```

Listing 1: Download: example-noauth.html

Portal page with authentication

Here is an example portal page requiring authentication.

```
<html>
<head>
        <title>Welcome to our portal</title>
</head> <body>
        <p>Welcome to our portal</p>
        <p>Enter your username and password and click Login to access the Internet</p> <form method="post"
        action="$PORTAL_ACTION$">
                <input name="auth_user" type="text">
                <input name="auth_pass" type="password">
                <input name="redirurl" type="hidden" value="$PORTAL_REDIRURL$">
                      <input name="zone" type="hidden" value="$PORTAL_ZONE$">
```

Listing 2: Download: example-auth.html

(continues on next page)

```
                <input name="accept" type="submit" value="Login">
        </form>
</body>
</html>
```

(continued from previous page)

Portal page with Vouchers

Here is an example portal page for use with vouchers.

```
<html>
<head>
            <title>Welcome to our portal</title>
</head> <body>
            <p>Welcome to our portal</p>
            <p>Enter your voucher code and click Login to access the Internet</p> <form method="post"
            action="$PORTAL_ACTION$">
                        <input name="auth_voucher" type="text">
                        <input name="redirurl" type="hidden" value="$PORTAL_REDIRURL$">
                                <input name="zone" type="hidden" value="$PORTAL_ZONE$">
                        <input name="accept" type="submit" value="Login">
            </form>
</body>
</html>
```

Listing 3: Download: example-voucher.html

Authentication error page contents

Using this control, optionally upload a custom HTML page to be displayed when authentication errors happen. An authentication error occurs when a user enters a bad username or password, or in the case of RADIUS authentication, potentially an unreachable RADIUS server.

By default, this error page is simply the login page again.

Logout page contents

The logout page is presented to the user after login and it triggers a popup window. The default code uses JavaScript to create the new window in the following way:

```
<html>
<head><title>Redirecting...</title></head> <body>
<span style="font-family: Tahoma, Verdana, Arial, Helvetica, sans-serif; font-size:␣
↪11px;">
<b>Redirecting to <a href="<?=$my_redirurl;?>"><?=$my_redirurl;?></a>...</b> </span>
<script type="text/javascript">
//<![CDATA[
LogoutWin = window.open('', 'Logout', 'toolbar=0,scrollbars=0,location=0,statusbar=0,
```

Listing 4: Download: example-logout.html

```
        ↪menubar=0,resizable=0,width=256,height=64');
```

(continues on next page)

```
if (LogoutWin) {
                LogoutWin.document.write('<html>');
                LogoutWin.document.write('<head><title>Logout</title></head>') ;
                LogoutWin.document.write('<body style="background-color:#435370">');
                LogoutWin.document.write('<div class="text-center" style="color: #ffffff;␣
↪font-family: Tahoma, Verdana, Arial, Helvetica, sans-serif; font-size: 11px;">') ;
                LogoutWin.document.write('<b>Click the button below to disconnect</b><p />');
                LogoutWin.document.write('<form method="POST" action="<?=$logouturl;?>">');
                LogoutWin.document.write('<input name="logout_id" type="hidden" value="<?=
↪$sessionid;?>" />');
                LogoutWin.document.write('<input name="zone" type="hidden" value="<?=$cpzone;?
↪>" />');
                LogoutWin.document.write('<input name="logout" type="submit" value="Logout" />
↪');
                LogoutWin.document.write('</form>');
                LogoutWin.document.write('</div></body>');
                LogoutWin.document.write('</html>');
                LogoutWin.document.close(); }

document.location.href="<?=$my_redirurl;?>"; //]]>
</script>
</body>
</html>
```

(continued from previous page)

Most browsers have pop-up blockers that will most likely stop that logout window from appearing, so investigate other possible means of creating a JavaScript pop-up using similar code.

20.4 MAC Address Control

The MACs tab defines actions for MAC addresses that can be either passed through the portal for this zone without requiring authentication, or blocked from reaching the portal.

To manage these MAC entries:

- Navigate to Services > Captive Portal
- Click ⧉ on the line for the Zone to edit
- Click the MACs tab
- Click + Add to add a new entry
- Fill in the form as follows:

  Action Defines the action to take on this entry:

    Pass Always allow traffic through from this MAC address without authentication.

    Block Always deny traffic from this MAC address

  MAC address The MAC address of the device to allow. The value must be colon-separated pairs of digits, such as 00:11:22:33:44:55.

  Description Some text describing the entry, if desired.

25.4. MAC Address Control

Bandwidth up/down The amount of bandwidth that this device may use, specified in Kilobits per second. Leave blank to not specify a limit.

- Click Save

From this page, an entry may be edited by clicking ⧉ on its row, or deleted by clicking 🗑.

20.5 Allowed IP Address

The Allowed IP Address tab works similarly to the MACs tab, except it checks IP addresses instead of MAC addresses. Traffic matching the specified IP address and the configured direction will always be allowed through the portal with no authentication in this zone.

IP Address The IP address of the device to always pass through the portal.

Description Some text describing the entry, if desired.

Direction The direction to allow traffic matching this IP address.

From Allow traffic sourced from this IP address through the portal, such as a local client IP address attempting to reach the Internet, or the IP address of a management client that must reach hosts on the portal network.

To Allow traffic with this IP address as a destination, such as a local web server IP address that must be reached via port forward, or a remote web server IP address which clients must always reach.

Both Allow traffic both to and from this IP address.

Bandwidth up/down The amount of bandwidth that this device may use. Leave blank to not specify a limit.

20.6 Allowed Hostnames

Allowed Hostnames work similarly to Allowed IP Address entries, except they are configured by hostname instead of IP address. A daemon periodically resolves the hostnames to IP address(es) and allows them through the portal without authentication in this zone.

The most common use of this feature is to make a "walled garden" style portal, where users are permitted to access a restricted set of sites without authenticating to the portal. This is also commonly used with the Pre-authentication Redirect URL if that page is hosted externally.

Note: Often sites will use many hostnames, content delivery networks, or ad servers as part of their content. In order to allow a site to load fully, all of these additional sites must be added to the list of allowed hostnames.

Direction The direction to allow traffic matching this hostname. In most typical use cases for allowing hostnames, the *To* or *Both* directions are the best fit.

To Allow traffic from local clients to a remote site matching this hostname as a destination without authentication. For example, a remote web server that must always be reachable by local clients, even when they are not logged in.

From Allow traffic sourced from this hostname through the portal, such as the hostname of a local client attempting to reach the Internet.

25.5. Allowed IP Address

Both Allow traffic both to and from this hostname.

Hostname The fully qualified domain name (FQDN) of the target host or site. The hostname must exist in DNS so that it can be resolved to an IP address.

Description Some text describing the entry, if desired.

Bandwidth up/down The amount of bandwidth that traffic to or from this hostname may use. Leave blank to not specify a limit.

20.7 Vouchers

Vouchers are single use codes used to gain Internet access through a Captive Portal. Each roll of vouchers is cryptographically generated and includes a set time limit. Vouchers are common in places where an organization wants authenticated, but time-limited, Internet access without needing to provide a username and password to users.

This type of configuration is prevalent in travel and hospitality locations such as coffee shops, hotels, and airports. Users enter a voucher code in the portal login form and are the portal grants access for the amount of time allowed by the voucher roll. Voucher rolls can be exported by the GUI as a CSV file, and some companies have even integrated the exported voucher lists into point of sale applications to print voucher codes on customer receipts.

Voucher time does not stop counting down if a user logs out; they allow access only from the start of the session until the duration of the voucher length has elapsed. During that time, the voucher can be re-used by the same or a different computer. If the voucher is used again by another computer, the previous session is stopped.

Vouchers require a public/private RSA key pair to generate and verify. A 32-bit set of keys is generated automatically by the firewall the first time the page loads. A custom key pair may be generated manually by the GUI as well. The maximum key length is 64 Bits. Using shorter keys will make the voucher codes shorter but eventually less secure.

20.7.1 Voucher Options

Voucher options are unique per Captive Portal Zone. To configure vouchers, navigate to the Vouchers tab when editing a Captive Portal Zone.

Voucher Rolls Voucher rolls are managed by this section of the page. The page lists information about each roll along with a link to add new rolls. No options appear here until after the other settings are active. See *Managing Voucher Rolls* for details.

Enable When checked, this portal zone allows vouchers as a method of authentication, and the page changes to display the other voucher options.

Voucher Keys The keys which generate and verify vouchers. Before vouchers are active on a zone, the GUI randomly generates new values each time the page loads. After saving voucher settings, the keys are present in the configuration and remain static from that point on.

> Warning: Do not change the keys or other bits after creating voucher rolls. If the values change, all current voucher rolls are invalid. Create new voucher rolls using the new settings after making changes.

Voucher Public Key This key is used by the portal to decrypt vouchers. Use the existing random key, or click Generate new keys to make a new public and private key pair. Users may generate keys elsewhere and paste the RSA public key (64 Bit or smaller) in PEM format here.

Voucher Private Key This key is used by the portal to generate voucher codes and does not need to be available if the vouchers are generated by another system. Use the existing random key or paste in an RSA private key (64 Bit or smaller) in PEM format here.

Character Set The character set defines which characters are valid for voucher text. The character set is case sensitive and should contain printable characters (numbers, lower case and upper case letters) that are hard to confuse with others. For example, avoid 0 (Digit zero), O (Letter O), l (Lowercase L), and 1 (Digit One). It cannot contain a space, double quote, or comma. A smaller character set will result in longer vouchers to ensure sufficient randomness.

Voucher Bits The following “bit” fields control how the vouchers themselves are generated by the portal. The best practice is to leave these values at their defaults, but they may be adjusted if necessary. The total of all bit fields *must* be less than the RSA key size. For example, the default values are 16, 10, and 5. The sum of these is 31, which is one less than the default RSA key size of 32.

# of Roll Bits Number of bits for the Roll ID. Set this larger to have a lot of rolls active at the same time. Can be from 1-31, the default value is 16.

# of Ticket Bits Number of bits for the Ticket ID. Set this larger if each roll will have a large number of vouchers. Can be from 1-16, the default value is 10.

# of Checksum Bits Reserves a range in each voucher to store a simple checksum over Roll bits and Ticket bits. Allowed range is 0-31, the default value is 5.

Magic Number The magic number is present in every voucher, and is verified by the portal during voucher check. The size of the magic number depends on the number of bits remaining after adding together the number of bits for the roll, ticket, and checksum. If there are no bits remaining for use by the magic number, then the portal does not use magic numbers.

Invalid Voucher Message This message is displayed by the portal to the user if they attempt to enter a voucher that does not exist or is not valid in any way except for being expired.

Expired Voucher Message This message is displayed to the user by the portal if they enter a voucher that was valid, but has expired.

20.7.2 Enable Vouchers

- Use the AZTCO-FW® WebGUI to navigate to Services > Captive Portal
- Click ⧉ on the line for the Zone to edit
- Ensure the Zone Authentication Method is set to *Use an Authentication backend*, change the value and save if necessary.
- Click the Vouchers tab

- Check Enable
- Fill in the form based on the options described in *Voucher Options*. In most cases, the options may remain at their default values.
- Click Save

With Vouchers enabled, the voucher management controls will be active in the GUI.

20.7.3 Managing Voucher Rolls

Vouchers are created by the portal in batches called Rolls. Each roll has specific settings that are unique to that roll. For example, a roll can have an 8-hour time limit and a separate roll can have a 12-hour time limit. Then users may be given voucher codes depending on which level of service they purchased and they will be limited to the amount of time corresponding to the voucher roll from which their code was picked.

Voucher Roll Options

Roll # The number of this roll. Each roll must have a unique number. This can be any number from 0 to 65535 with the default number of Roll Bits.

Minutes per Ticket Defines how long the voucher lasts, in minutes. The voucher time starts counting down the moment the voucher is used, and does not stop, so plan the voucher length accordingly.

Because this is defined in minutes, ensure the correct length is used, e.g. 1440 minutes is 24 hours.

Count Defines the number of vouchers in this roll. The value can be from 0 to 1023 with the default number of Ticket Bits.

Note: If the count on an existing roll is changed, it will invalidate all other vouchers on the roll.

Comment A description of the roll for reference, such as 2 hour vouchers for coffee purchases.

Creating Voucher Rolls

To create a voucher roll:

- Use the AZTCO-FW® WebGUI to navigate to Services > Captive Portal
- Click ⧉ on the line for the Zone to edit
- Click the Vouchers tab
- Click ✚ Add under the roll list
- Fill in the options as described in *Voucher Roll Options*
- Click Save

The new roll is available for immediate use by clients.

Editing Existing Rolls

To edit an existing voucher roll, click ⧉ at the end of its row, but be careful when making changes. Changing the Roll number or Count will invalidate the current vouchers on the roll.

Removing Voucher Rolls

To remove rolls of vouchers, click 🗑 at the end of their row. When a roll is removed, *all* of the vouchers in that roll become invalid. Do not remove a roll unless it has been completely used, compromised, or otherwise unnecessary.

Exporting/Downloading Voucher Rolls

Click X to download a file containing the vouchers in the specified roll. This action downloads a .csv (Comma Separated Value) spreadsheet containing all voucher codes for this roll.

Nearly any spreadsheet editor can open this file, such as LibreOffice Calc, Google Docs, or Excel. Programs such as those can print vouchers, feed them into a POS system, and so on.

Using Vouchers on A Portal Page

The portal page must submit voucher codes via the auth_voucher form field. See *Portal page with Vouchers* for an example.

Viewing Active Vouchers

To view the list of currently active vouchers and their timers, navigate to Status > Captive Portal, on the Active Vouchers tab for a zone, as seen in Figure *Active Vouchers*.

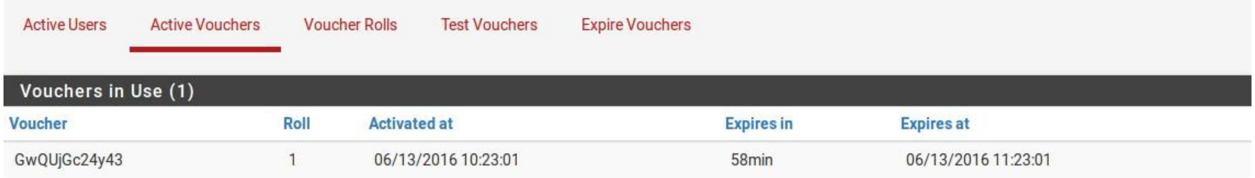

Vouchers in Use (1)

Voucher	Roll	Activated at	Expires in	Expires at
GwQUjGc24y43	1	06/13/2016 10:23:01	58min	06/13/2016 11:23:01

Fig. 1: Active Vouchers

Viewing Voucher Roll Utilization

To view a list of voucher rolls and usage counts, navigate to Status > Captive Portal, on the Voucher Rolls tab for a zone, as in Figure *Vouchers Roll Usage*

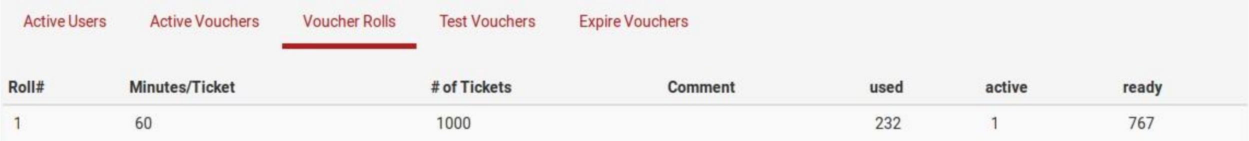

Roll#	Minutes/Ticket	# of Tickets	Comment	used	active	ready
1	60	1000		232	1	767

Fig. 2: Vouchers Roll Usage

Testing Vouchers

To test the validity of a voucher code, enter it at Status > Captive Portal, on the Test Vouchers tab for a zone. Upon submission, the page will display if a code is valid or not, and if it is valid, it will show the voucher time limit, as seen in Figure *Testing Vouchers*. Testing a voucher does not count it as used or expired, it is still free to be used at a later time by a client.

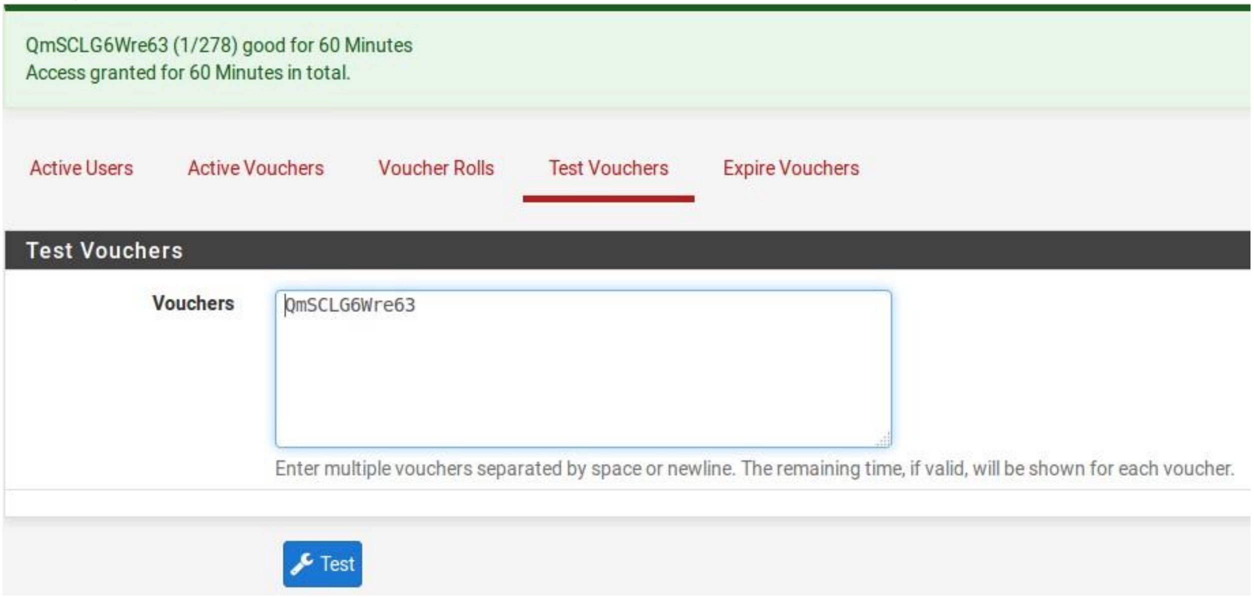
QmSCLG6Wre63 (1/278) good for 60 Minutes
Access granted for 60 Minutes in total.

Test Vouchers

Vouchers: QmSCLG6Wre63

Enter multiple vouchers separated by space or newline. The remaining time, if valid, will be shown for each voucher.

Fig. 3: Testing Vouchers

Expiring Vouchers

To invalidate vouchers, during or before use, enter them at Status > Captive Portal, on the Expire Vouchers tab for a zone. After submitting, any voucher listed in the form will no longer be allowed by the portal. Active vouchers entered on this page are also immediately expired.

The page can expire any number of vouchers in a batch. Enter vouchers separated by newlines.

20.7.4 Synchronizing Vouchers

At the bottom of the Vouchers tab there are options to synchronize vouchers to another HA node. This works similarly to the XML-RPC configuration synchronization found in high availability setups (See *AZTCO-FW XML-RPC Config Sync Overview*). When active, this function copies the voucher rolls to the target node and also pushes information about active vouchers to the target node as vouchers are consumed by users.

Synchronize Voucher Database IP The target IP address or hostname of the other node for voucher synchronization.

Voucher sync port The port on the target node where the GUI is listening (Typically 443).

Voucher sync username The username for synchronization access (Must have appropriate privileges).

Voucher sync password The GUI password for the target node.

Unlike the configuration synchronization in a high availability cluster, this synchronization is configured on both the primary and secondary nodes. This ensures that vouchers used on the secondary node while it is active are also sent back to the primary when it returns to an active state. Unlike configuration synchronization, this does not create a loop.

20.8 File Manager

The File Manager tab in a Captive Portal zone is used to upload files that can then be utilized inside a captive portal page, such as style sheets, image files, PHP or JavaScript files.

The total size limit for all files in a zone is 1 MB.

20.8.1 File Name Conventions

When a file is uploaded using the File Manager, the file name will automatically be prefixed with captiveportal-. For example, if logo.png is uploaded it will become captiveportal-logo.png. If a file already has that prefix in its name, the name is not changed.

These files will be made available in the root directory of the captive portal server for this zone. The files may be referenced directly from the portal page HTML code using relative paths.

Example: An image with the name captiveportal-logo.jpg was uploaded using the file manager, It can then be included in the portal page as follows:

```
<img src="captiveportal-logo.jpg" />
```

PHP scripts may be uploaded as well, but they may need to be passed extra parameters to work as desired, for example:

```
<a href="/captiveportal-aup.php?zone=$PORTAL_ZONE$&redirurl=$PORTAL_REDIRURL$">
        Acceptable usage policy
</a>
```

20.8.2 Managing Files

To upload files:

- Navigate to Services > Captive Portal
- Edit the zone where the files will be uploaded
- Click the File Manager tab
- Click +
- Click Browse
- Locate and select the file to upload
- Click ⬆ Upload

The file will be transferred to the firewall and stored in the configuration.

To delete files:

- Navigate to Services > Captive Portal
- Edit the zone where the file to delete is located

25.8. File Manager

- Click the File Manager tab
- Click 🗑 next to the file to remove
- Click OK to confirm the delete action

The file will be removed from the portal configuration and will no longer be available for use in portal pages.

See also:

- *Captive Portal Status*
- *Captive Portal Authentication Logs*
- *Troubleshooting Captive Portal*

Captive Portal in AZTCO-FW® software forces users on an interface to authenticate before granting access to the Internet. Where possible, the firewall automatically presents a login web page in which the user must enter credentials such as a username/password, a voucher code, or a simple click-through agreement.

This feature is commonly used throughout the hospitality industry (Hotels, restaurants, airports, and more) as well as in corporate and even home environments. It is primarily used for wireless hot spots or for additional authentication before allowing access to internal networks from wireless clients.

Captive Portal is configured under Services > Captive Portal.

20.9 Limitations

The Captive Portal implementation in AZTCO-FW does have some limitations. This section covers those, and the common ways of working around them where possible.

20.9.1 Does not yet support IPv6

Currently, Captive Portal does not support IPv6.

20.9.2 Not capable of reverse portal

A reverse portal, requiring authentication for traffic coming into a local network from the Internet, is not possible.

CHAPTER

TWENTYONE

HIGH AVAILABILITY

26.1 pfsync Overview

pfsync enables the synchronization of the firewall state table between cluster nodes. Changes to the state table on the primary are sent to the secondary firewall(s) over the Sync interface, and vice versa. When pfsync is active and properly configured, all nodes will have knowledge of each connection flowing through the cluster. If the master node fails, the backup node will take over and clients will not notice the transition since both nodes knew about the connection beforehand.

pfsync uses multicast by default, though an IP address can be defined to force unicast updates for environments with only two firewalls where multicast traffic will not function properly. Any active interface can be used for sending pfsync updates, however utilizing a dedicated interface is better for security and performance. pfsync does not support any method of authentication, so if anything other than a dedicated interface is used, it is possible for any user with local network access to insert states into the state table. In low throughput environments that aren't security paranoid, use of the LAN interface for this purpose is acceptable. Bandwidth required for this state synchronization will vary significantly from one environment to another, but could be as high as 10% of the throughput traversing the firewall depending on the rate of state insertions and deletions in a network.

Failover can still operate without pfsync, but it will not be seamless. Without pfsync if a node fails and another takes over, user connections would be dropped. Users may immediately reconnect through the other node, but they would be disrupted during the transition. Depending on the usage in a particular environment, this may go unnoticed or it could be a significant, but brief, outage.

When pfsync is in use, pfsync settings *must be enabled* on all nodes participating in state synchronization, including secondary nodes, or it will not function properly.

21.1.1 pfsync and Firewall Rules

Traffic for pfsync must be explicitly passed on the Sync interface. The rule must pass the *pfsync* protocol from a source of the *Sync network* to *any* destination. A rule passing all traffic of any protocol would also allow the required traffic, but a more specific rule is more secure.

21.1.2 pfsync and Physical Interfaces

States in AZTCO-FW® are bound to specific operating system Interfaces. For example, if WAN is *em0*, then a state on WAN would be tied to *em0*. If the cluster nodes have identical hardware and interface assignments then this works as expected. In cases when different hardware is used, this can be a problem. If WAN on one node is *em0* but on another node it is *igb0*, the states will not match and they will not be treated the same.

It is always preferable to have identical hardware, but in cases where this is impractical there is a workaround: Adding interfaces to a LAGG will abstract the actual underlying physical interface so in the above example, WAN would be *lagg0* on both and states would be bound to *lagg0*, even though *lagg0* on one node contains *em0* and it contains *igb0* on the other node.

21.1.3 pfsync and Upgrades

Normally AZTCO-FW would allow firewall upgrades without any network disruption. Unfortunately, this isn't always the case with upgrades as the pfsync protocol can change to accommodate additional functionality. Always check the upgrade guide linked in all release announcements before upgrading to see if there are any special considerations for CARP users.

21.2 AZTCO-FW XML-RPC Config Sync Overview

To make the job of maintaining practically identical AZTCO-FW® software nodes easier, configuration synchronization is possible using XML-RPC. When XML-RPC Synchronization is enabled, settings from supported areas are copied to the secondary and activated after each configuration change. XMLRPC Synchronization is optional, but maintaining a cluster is a lot more work without it.

Some areas cannot be synchronized, such as the Interface configuration, but many other areas can: Firewall rules, aliases, users, certificates, VPNs, DHCP, routes, gateways, and more. As a general rule, items specific to hardware or a particular installation, such as Interfaces or values under System > General or System > Advanced do not synchronize. The list of supported areas can vary depending on the version of AZTCO-FW in use. For a list of areas that will synchronize, see the checkbox items on System > High Avail Sync in the XMLRPC section. Most packages will not synchronize but some contain their own synchronization settings. Consult package documentation for more details.

Configuration synchronization should use the Sync interface, or if there is no dedicated Sync interface, use the same interface configured for pfsync.

In a two-node cluster the XML-RPC settings must *only* be enabled on the primary node, the secondary node must have these settings *disabled*.

For XML-RPC to function, both nodes must meet the following requirements:

- The GUI must be running on the same port and protocol, for example: HTTPS on port 443, which is the default setting.
- The interfaces must be assigned identically on both nodes, for example: wan=WAN, lan=LAN, opt1=Sync, opt2=DMZ. Check the config.xml contents directly to ensure a match.

•

26.2. AZTCO-FW XML-RPC Config Sync Overview

> Warning: If the interfaces do not match up exactly, firewall rules and other configuration items will appear to synchronize to the wrong interface on the secondary node. Additionally, this can also lead to failures in DHCP failover.

- The sync user must either be admin or an account with the *System - HA node sync* privilege.

> Note: If XML-RPC will synchronize users, create the sync user on the secondary manually first, as well as on the primary. The redundant copy on the secondary will be removed during the first successful synchronization, but the initial synchronization cannot succeed without it.

21.3 Verifying Failover Functionality

Since using HA is about high availability, thorough testing before placing a cluster into production is a must. The most important part of that testing is making sure that the HA peers will failover gracefully during system outages.

If any actions in this section do not work as expected, see *Troubleshooting High Availability*.

21.3.1 Check CARP status

On both systems, navigate to Status > CARP (failover). If everything is working correctly, the primary will show ↻ MASTER for the status of all CARP VIPs and the secondary will show ⏸ BACKUP.

If either instead shows DISABLED, click the Enable CARP button and then refresh the page.

If an interface shows ⊗ INIT, it means the interface containing the CARP VIP does not have a link. Connect the interface to a switch, or at least to the other node. If the interface will not be used for some time, remove the CARP VIP from the interface as this will interfere with normal CARP operation.

21.3.2 Check Configuration Replication

Navigate to key locations on the secondary node, such as Firewall > Rules and Firewall > NAT and ensure that rules created only on the primary node are being replicated to the secondary node.

If the example earlier in this chapter was followed, the “temp” firewall rule on the pfsync interface would be replaced by the rule from the primary.

21.3.3 Check DHCP Failover Status

If DHCP failover was configured, its status can be checked at Status > DHCP Leases. A new section will appear at the top of the page containing the status of the DHCP Failover pool, as in Figure *DHCP Failover Pool Status*.

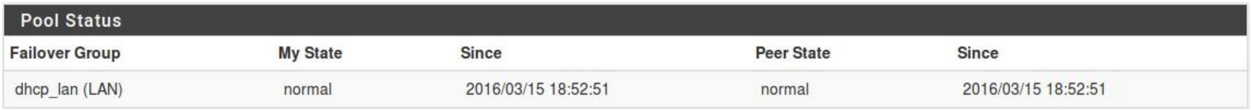

Pool Status

Failover Group	My State	Since	Peer State	Since
dhcp_lan (LAN)	normal	2016/03/15 18:52:51	normal	2016/03/15 18:52:51

Fig. 1: DHCP Failover Pool Status

26.3. Verifying Failover Functionality

21.3.4 Test CARP Failover

Now for the real failover test. Before starting, make sure that a local client behind the CARP pair on LAN can connect to the Internet with both AZTCO-FW® firewalls online and running. Once that is confirmed to work, it is an excellent time to make a backup.

For the actual test, unplug the primary node from the network or shut it down temporarily. The client will be able to keep loading content from the Internet through the secondary node. Check Status > CARP (failover) again on the backup and it will now report that it is MASTER for the LAN and WAN CARP VIPs.

Now bring the primary node back online and it will regain its role as MASTER, and the backup system will demote itself to BACKUP once again. At any point during this process, Internet connectivity will still work properly.

Test the HA pair in as many failure scenarios as possible. Additional tests include:

- Unplug the WAN or LAN cable
- Pull the power plug of the primary
- Disable CARP on the primary using both the temporary disable feature and maintenance mode
- Test with each system individually (power off secondary, then power back on and shut down the primary)
- Download a file or try streaming audio/video during the failover
- Run a continuous ICMP echo request (ping) to an Internet host during the failover

21.4 Layer 2 Redundancy

The diagrams earlier in this chapter did not describe layer 2 (switch) redundancy, to avoid throwing too many concepts at readers simultaneously. This section covers the layer 2 design elements to be considered when planning a redundant network. This chapter assumes a two system deployment, though this scales to as many installations as required.

If both redundant AZTCO-FW® firewalls are plugged into the same switch on any interface, that switch becomes a single point of failure. To avoid this single point of failure, the best choice is to deploy two switches for each interface (other than the dedicated pfsync interface).

Example HA Network Diagram is network-centric, not showing the switch infrastructure. The Figure *Diagram of HA with Redundant Switches* illustrates how that environment looks with a redundant switch infrastructure.

21.4.1 Switch Configuration

When using multiple switches, the switches should be interconnected. As long as there is a single connection between the two switches, and no bridge on either of the firewalls, this is safe with any type of switch. Where using bridging, or where multiple interconnections exist between the switches, care must be taken to avoid layer 2 loops. A managed switch would be required which is capable of using Spanning Tree Protocol (STP) to detect and block ports that would otherwise create switch loops. When using STP, if an active link dies, e.g. switch failure, then a backup link can automatically be brought up in its place.

AZTCO-FW also has support for *lagg(4)* link aggregation and link failover interfaces which allows multiple network interfaces to be plugged into one or more switches for increased fault tolerance. See *LAGG (Link Aggregation)* for more information on configuring link aggregation.

26.4. Layer 2 Redundancy

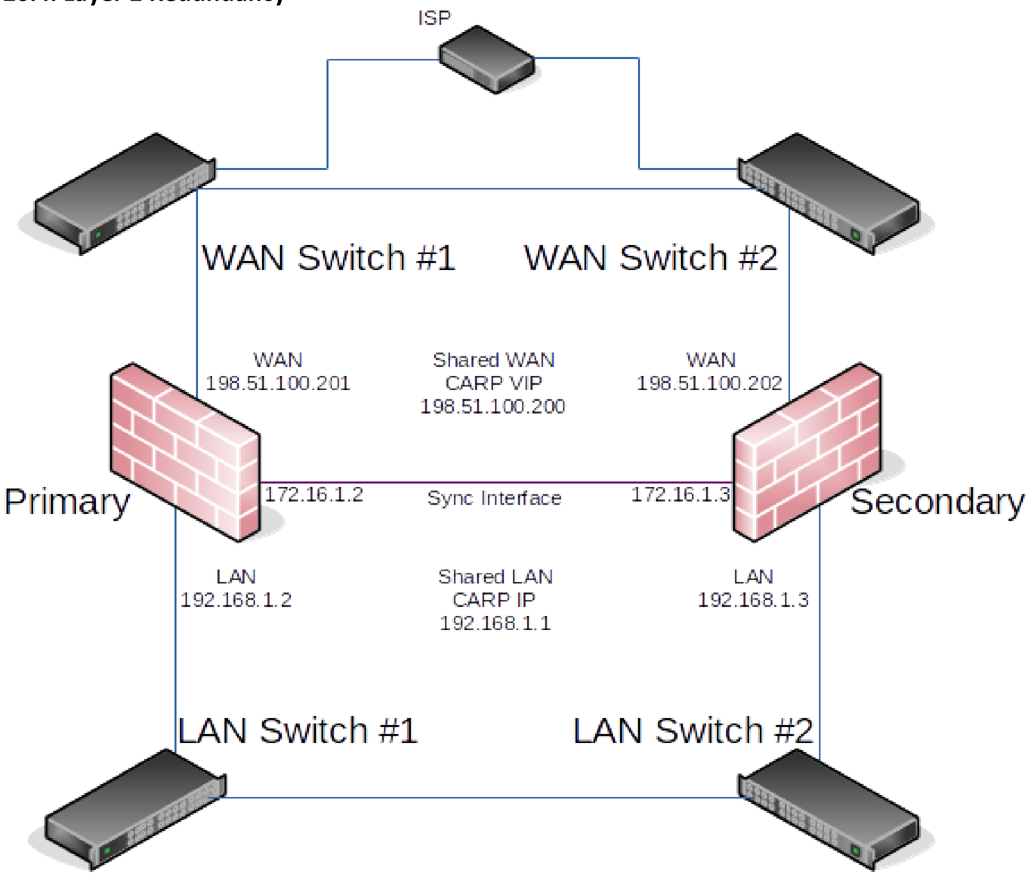

Fig. 2: Diagram of HA with Redundant Switches

21.4.2 Host Redundancy

It is more difficult to obtain host redundancy for critical systems inside the firewall. Each system could have two network cards and a connection to each group of switches using Link Aggregation Control Protocol (LACP) or similar vendor-specific functionality. Servers could also have multiple network connections, and depending on the OS it may be possible to run CARP or a similar protocol on a set of servers so that they would be redundant as well. Providing host redundancy is more specific to the capabilities of the switches and server operating systems, which is outside the scope of this documentation.

21.4.3 Other Single Points of Failure

When trying to design a fully redundant network, there are many single points of failure that sometimes get missed. Depending on the level of uptime to achieve, there are more and more things to consider than a simple switch failure. Here are a few more examples for redundancy on a wider scale:

- Supply isolated power for each redundant segment.
  - Use separate breakers for redundant systems.
  - Use multiple UPS banks/generators.
  - Use multiple power providers, entering opposite sides of the building where possible.
- Even a Multi-WAN configuration is no guarantee of Internet uptime.
  - Use multiple Internet connection technologies (DSL, Cable, Fiber, Wireless).
  - If any two carriers use the same pole/tunnel/path, they could both be knocked out at the same time.
- Have backup cooling, redundant chillers or a portable/emergency air conditioner.
- Consider placing the second set of redundant equipment in another room, another floor, or another building.
- Have a duplicate setup in another part of town or another city.
- I hear hosting is cheap on Mars, but the latency is killer.

21.5 High Availability with Bridging

High availability is not currently compatible with bridging in a native capacity that is considered reliable or worthy of production use. It requires significant manual intervention. The details of the process can be found in *High Availability*.

21.6 Using IP Aliases to Reduce Heartbeat Traffic

If there are a large number of CARP VIPs on a segment, this can lead to a lot of multicast traffic. One heartbeat per second is sent per CARP VIP. To reduce this traffic, additional VIPs may be “stacked” on top of one CARP VIP on an interface. First, pick one CARP VIP to be the “main” VIP for the interface. Then, change the other CARP VIPs in that same subnet to be an *IP Alias* type VIP, with the “main” CARP VIP interface selected to be their Interface on the VIP configuration.

This not only reduces the heartbeats that will be seen on a given segment, but it also causes all of the IP alias VIPs to change status along with the "main" CARP VIP, reducing the likelihood that a layer 2 issue will cause individual CARP VIPs to not fail over as expected.

IP Alias VIPs do not normally synchronize via XML-RPC configuration synchronization, however, IP alias VIPs set to use CARP interfaces in this manner will synchronize.

26.5. High Availability with Bridging

AZTCO-FW® software is one of very few open source solutions offering enterprise- class high availability capabilities with stateful failover, allowing the elimination of the firewall as a single point of failure. High Availability is achieved through a combination of features:

- CARP for IP address redundancy
- XMLRPC for configuration synchronization
- pfsync for state table synchronization

With this configuration, nodes act as an "active/passive" cluster with the primary node working as the master node and the secondary node in a backup role, taking over as needed if the primary node fails.

Though often erroneously called a "CARP Cluster", two or more redundant firewalls are more aptly titled a "High Availability Cluster" or "HA Cluster", since CARP is only one of several technologies used to achieve High Availability with AZTCO-FW, and in the future CARP could be swapped for a different redundancy protocol.

One interface on each cluster node will be dedicated for synchronization tasks. This is typically referred to as the "Sync" interface, and it is used for configuration synchronization and pfsync state synchronization. Any available interface may be used.

Note: Some call this the "CARP" interface but that is incorrect and very misleading. CARP heartbeats happen on each interface with a CARP VIP; CARP traffic and failover actions do not utilize the Sync interface.

The most common High Availability cluster configuration includes only two nodes. It is possible to have more nodes in a cluster, but they do not provide a significant advantage.

It is important to distinguish between the three functions (IP address redundancy, configuration synchronization, and state table synchronization), because they happen in different places. Configuration synchronization and state synchronization happen on the sync interface, directly communicating between firewall units. CARP heartbeats are sent on each interface with a CARP VIP. Failover signaling does not happen on the sync interface, but rather it happens on every CARP-enabled interface.

See also:

AZTCO-FW Hangouts on Youtube which contains the June 2015 Hangout also covering High Availability.

26.7 CARP Overview

Common Address Redundancy Protocol (CARP) was created by OpenBSD developers as a free, open redundancy solution for sharing IP addresses among a group of network devices. Similar solutions already existed, primarily the IETF standard for Virtual Router Redundancy Protocol (VRRP). However Cisco claims VRRP is covered by its patent on their Hot Standby Router Protocol (HSRP), and told the OpenBSD developers that it would enforce its patent. Hence, the OpenBSD developers created a new free, open protocol to accomplish essentially the same result without infringing on Cisco's patent. CARP became available in October 2003 in OpenBSD, and was later added to FreeBSD as well.

A CARP type Virtual IP address (VIP) is shared between nodes of a cluster. One node is master and receives traffic for the IP address, and the other nodes maintain backup status and monitor for heartbeats to see if they need to assume the master role if the previous master fails. Since only one member of the cluster at a time is using the IP address, there is no IP address conflict for CARP VIPs.

In order for failover to work properly it is important that inbound traffic coming to the cluster, such as routed upstream traffic, VPNs, NAT, local client gateway, DNS requests, etc., be sent to a CARP VIP and for outgoing traffic such as Outbound NAT to be sent from a CARP VIP. If traffic is addressed to a node directly and not a CARP VIP, then that traffic will not be picked up by other nodes.

CARP works similar to VRRP and HSRP, and may even conflict in some cases. Heartbeats are sent out on each interface containing a CARP VIP, one heartbeat per VIP per interface. At the default values for skew and base, a VIP sends out heartbeats about once per second. The skew determines which node is master at a given point in time. Whichever node transmits heartbeats the fastest assumes the master role. A higher skew value causes heartbeats to be transmitted with more delay, so a node with a lower skew will be the master unless a network or other issue causes the heartbeats to be delayed or lost.

Note: Never access the firewall GUI, SSH, or other management mechanism using a CARP VIP. For management purposes, only use the actual IP address on the interface of each separate node and not the VIP. Otherwise it cannot be determined beforehand which unit is being accessed.

21.7.1 IP Address Requirements for CARP

A High Availability cluster using CARP needs three IP addresses in each subnet along with a separate unused subnet for the Sync interface. For WANs, this means that a /29 subnet or larger is required for an optimal configuration. One IP address is used by each node, plus a shared CARP VIP address for failover. The synchronization interface only requires one IP address per node.

It is technically possible to configure an interface with a CARP VIP as the only IP address in a given subnet, but it is not generally recommended. When used on a WAN, this type of configuration will only allow communication from the primary node to the WAN, which greatly complicates tasks such as updates, package installations, gateway monitoring, or anything that requires external connectivity from the secondary node. It can be a better fit for an internal interface, however internal interfaces do not typically suffer from the same IP address limitations as a WAN, so it is still preferable to configure IP addresses on all nodes.

21.7.2 Switch/Layer 2 Concerns

CARP heartbeats utilize multicast and may require special handling on the switches involved with the cluster. Some switches filter, rate limit, or otherwise interfere with multicast in ways that can cause CARP to fail. Also, some switches employ port security methods which may not work properly with CARP.

At a minimum, the switch must:

- Allow Multicast traffic to be sent and received without interference on ports using CARP VIPs.
- Allow traffic to be sent and received using multiple MAC addresses.
- Allow the CARP VIP MAC address to move between ports.

Nearly all problems with CARP failing to properly reflect the expected status are failures of the switch or other layer 2 issues, so be sure the switches are properly configured before continuing.

CHAPTER

TWENTYTWO

SYSTEM MONITORING

The data and information that AZTCO-FW® software collects and displays is every bit as important as the services it provides. Sometimes it seems that commercial routers go out of their way to hide as much information as possible from users, but AZTCO-FW can provide almost as much information as anyone could ever want (and then some).

This chapter contains a variety of methods for finding information about the firewall status, logs, traffic, hardware, and so on.

22.1 Status

These articles cover various ways to check the status of services or features of the firewall, or the firewall itself.

22.1.1 Dashboard

The main page of the firewall is the Dashboard. The Dashboard page provides a wealth of information that can be seen at a glance, contained in configurable widgets. These widgets can be added or removed, and dragged around into different positions.

Managing Widgets

Each widget follows some basic conventions for controlling its position, size, settings, and so on, the mechanics of which are covered here in this section.

Adding and Removing Widgets

To start adding widgets, click the + button in the Dashboard controls area of the breadcrumb bar to display the list of available widgets. See *Dashboard Controls in the Breadcrumb Bar*.

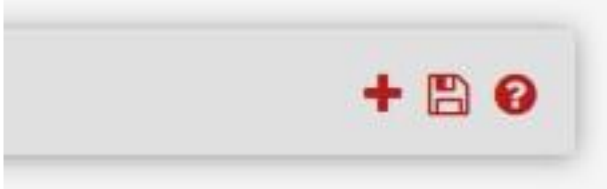

Fig. 1: Dashboard Controls in the Breadcrumb Bar

Inside the Available Widgets panel, click on the name of a widget to add it to the Dashboard (See *Available Widgets List*). The dashboard will reload with the new widget displayed in one of its columns.

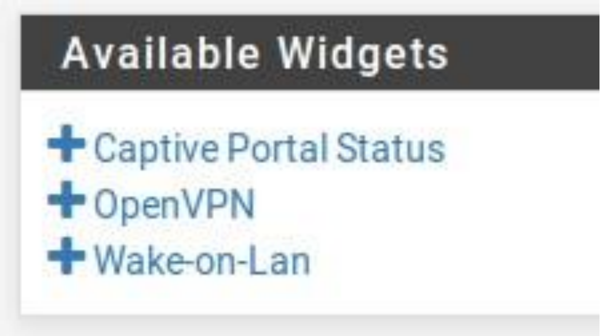

Fig. 2: Available Widgets List

To close and remove a widget from the Dashboard, click the ⊗ button in its title bar, as seen in Figure *Widget Title Bar*, then click 💾 in the dashboard controls.

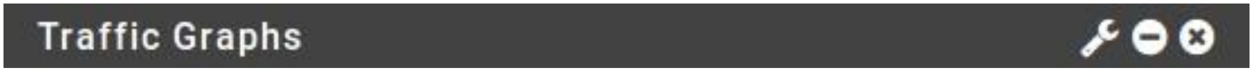

Fig. 3: Widget Title Bar

Rearranging Widgets

Widgets can be rearranged and moved between columns. To move a widget, click and drag its title bar (Figure *Widget Title Bar*), move the mouse to the desired position, and then release. As the widget is moved it will "snap" into its new position, so the new location may be previewed before releasing the mouse button. After positioning a widget, click

💾 in the dashboard controls (*Dashboard Controls in the Breadcrumb Bar*).

Minimizing Widgets

To minimize a widget so it hides its content and only shows up as its title bar, click ⊖ thebutton in its title bar, ⊕

as seen in Figure *Widget Title Bar*. To restore the widget to its normal display, click thebutton. After changing the widget status, click 💾 in the dashboard controls (*Dashboard Controls in the Breadcrumb Bar*).

Changing Widget Settings

Some widgets have customizable settings that control how their content is displayed or updated. If a widget has settings, the 🔧 button will show up in its title bar as seen in Figure *Widget Title Bar*. Click that button and the settings for the widget will appear. Once the settings have been adjusted, click the Save button inside of the widget settings panel.

Available Widgets

Each widget contains a specific set of data, type of information, graph, etc. Each of the currently available widgets will be covered in this section, along with their settings (if any). These are listed in alphabetical order.

Captive Portal Status

This widget shows the current list of online captive portal users, including their IP address, MAC address, and username.

CARP Status

The CARP Status widget displays a list of all CARP type Virtual IP addresses, along with their status as either MASTER or BACKUP.

Dynamic DNS

The Dynamic DNS widget displays a list of all configured Dynamic DNS hostnames, their current address, and status.

Gateways

The Gateways widget lists all of the system gateways along with their current status. The status information consists of the gateway IP address, Round Trip Time (RTT) also known as delay or latency, the amount of packet loss, and the status (Online, Warning, Down, or Gathering Data). The widgets is updated every few seconds via AJAX.

Gmirror Status

This widget will show the status of a gmirror RAID array on the system, if one is configured. The widget will show if the array is online/OK (Complete), rebuilding, or degraded.

Installed Packages

The Installed Packages widget lists all of the packages installed on the system, along with some basic information about them such as the installed version and whether or not an update is available.

When a package has an update available, ⊕ is displayed next to the version number. Packages may be updated from this widget by clicking the ↻ button at the end of a package’s row.

Packages may also be reinstalled by clicking ⇄ or removed by clicking 🗑.

Interface Statistics

This widget shows a grid, with each interface on the system shown in its own column. Various interface statistics are shown in each row, including packet, byte, and error counts.

Interfaces

The Interfaces widget differs from the Interface Statistics widget in that it displays general information about the interface rather than counters. The Interfaces widget shows the type and name of each interface, IPv4 address, IPv6 address, the interface link status (up or down), as well as the link speed when available.

IPsec

The IPsec widget has three tabs: The first tab, Overview, is a count of active and inactive tunnels. The second tab, Tunnel Status, lists each configured IPsec tunnel and whether that tunnel is up or down. The last tab, Mobile, shows online remote access IPsec VPN users, such as those using IKEv2 or Xauth.

Load Balancer Status

This widget displays a compact view of the server Load Balancing setup. Each row shows the status for one virtual server. The Server column shows the virtual server name, status, and IP address with port where the virtual server is accepting connections. The Pool column shows the individual pool servers and their status, with an uptime percentage. The Description column shows the text description from the virtual server.

Firewall Logs

The Firewall Logs widget provides an AJAX-updating view of the firewall log. The number of rows shown by the widget is configurable. As with the normal firewall log view, clicking the action icon next to the log entry will show a window displaying which rule caused the log entry. Clicking the source or destination IP address will copy that value to Diagnostics > DNS where the address can be resolved.

NTP Status

The NTP Status widget shows the current NTP synchronization source and the server time from that source.

OpenVPN

The OpenVPN widget displays the status of each configured OpenVPN instance, for both servers and clients. The status of each instance is shown, but the style and type of information shown varies depending on the type of OpenVPN connection. For example, SSL/TLS based servers show a list of all connected clients. For static key clients and servers, an up/down status is displayed. In each case it displays the IP address of the connecting client with the name and time of the connection.

Picture

The Picture widget, as the name implies, displays a picture chosen by the user. This can either be used functionally, for a network diagram or similar, or it can be for style, displaying a company logo or other image.

To add an image:

- Click 🔧 on the Picture widget title bar
- Click Browse to locate the picture to upload
- Click Upload to upload the picture

The size of the picture will adjust to fit the area of the widget, which can vary depending on the size of the browser and platform.

RSS

The RSS (RDFSite Summary, or as it’s often called, Really Simple Syndication) widget will display an arbitrary RSS feed. By default, it shows the AZTCO-FW® blog RSS feed. Some people choose to show internal company RSS feeds or security site RSS feeds, but it can load any RSS feed.

In addition to defining the RSS feeds to display, the number of stories and size of displayed content are also configurable.

Services Status

This widget provides the same view and control of services that appears under Status > Services. Each service is listed along with its description, status (Running, Stopped), and start/restart/stop controls.

SMART Status

If S.M.A.R.T. is enabled on a drive in the firewall, this widget will show a brief status of the drive integrity as reported by S.M.A.R.T.

System Information

This widget is the main widget, displaying a wide array of information about the running system. The information displayed includes:

Name The configured fully qualified hostname of the firewall.

Version The current running version of AZTCO-FW on the firewall. The version, architecture, and build time are displayed at the top. Under the build time, the underlying version of FreeBSD is shown.

Under those items is the result of an automatic update check for a more recent version of AZTCO-FW software. This automatic update check can be disabled in the update settings.

CPU Type The displayed CPU type is the version string for the processor, such as “Intel(R) Atom(TM) CPU C2758 @ 2.40GHz”. The CPU count and package/core layout is also displayed.

If powerd is active and the CPU frequency has been lowered, then the current frequency is shown along size the maximum frequency.

Hardware crypto If a known hardware cryptographic accelerator has been detected, it will be displayed here.

Uptime This is the time since the firewall was last rebooted.

Current date/time The current date and time of the firewall, including the time zone. This is useful for comparing the log entries, especially when the time zone on the firewall is different from where the user resides.

DNS Server(s) Lists all of the configured DNS Servers on the firewall.

Last config change The date of the last configuration change on the firewall.

State table size Shows a graphical and numerical representation of active states and the maximum possible states as configured on the firewall. Underneath the state counts is a link to view the contents of the state table.

MBUF usage Shows the number of network memory buffer clusters in use, and the maximum the system has available. These network memory buffers are used for network operations, among other tasks. If the number is close to maximum or at the maximum, increase the number of available mbufs as described in *Hardware Tuning and Troubleshooting*.

Load Average A count of how many active processes are running on the firewall during the last 5, 10, and 15 minutes. This is typically 0.00 on an idle or lightly loaded system.

CPU usage A bar chart and percentage of CPU time in use by the firewall. Note that viewing the dashboard will increase the CPU usage a bit, depending on the platform. On slower platforms this is likely to read significantly higher than it would be otherwise.

Memory usage The current amount of RAM in use by the system. Note that unused RAM is often allocated for caching and other tasks so it is not wasted or idle, so this number may show higher than expected even if it is operating normally.

Swap usage The amount of swap space in use by the system. If the system runs out of physical RAM, and there is swap space available, lesser used pages of memory will be paged out to the swap file on the hard drive. This indicator only shows when the system has swap space configured, which will only be on full installs.

Disk usage The amount of space used on the boot disk or storage media. The type and location of mounted filesystems are shown, including memory disks when present.

Thermal Sensors

The Thermal Sensors widget displays the temperature from supported sensors when present. For many popular Intel and AMD-based chips, the sensors may be activated by choosing the appropriate sensor type under System > Advanced on the Miscellaneous tab under Thermal Sensors

A bar is displayed for each sensor, which typically corresponds to each CPU core. The warning and critical thresholds may be configured in the widget settings.

Traffic Graphs

The Traffic Graphs widget contains a live SVG graph for the traffic on each interface. The interfaces displayed are configurable in the widget settings. The default refresh rate of the graphs is once every 10 seconds, but that may also be adjusted in the settings for this widget. The graphs are drawn the same way as those found under Status Traffic Graph.

Wake On LAN

The Wake on LAN widget shows all of the WOL entries configured under Services Wake on LAN, and offers a quick means to send the magic packet to each system in order to wake it up. The current status of a system is also shown. To

wake up a system, click ⏻ next to its entry.

22.1.2 Interface Status

The status of network interfaces may be viewed at Status > Interfaces. In the first part of Figure *Interface Status*, a DHCP WAN connection has been made and the IPv4 and IPv6 address, DNS, etc have been obtained automatically. The MAC address, media type, in/out packets, errors, and collisions for the network interface are all visible. Dynamic connection types like PPPoE and PPTP have a Disconnect button when connected and a Connect button when offline. Interfaces obtaining an IP address from DHCP have a Release button when there is an active lease, and a Renew button when there is not.

In the lower part of the image, the LAN connection is visible. Since this is a normal interface with a static IP address, only the usual set of items are shown.

If an interface status indicates “no carrier” then it typically means that the cable is not plugged in or the device on the other end is malfunctioning in some way. If any errors are shown, they are typically physical in nature: cabling or port errors. The most common suspect is cables, and they are easy and cheap to replace. In some circumstances errors and collisions may appear due to a link speed or duplex mismatch. See *Speed and Duplex* for more about setting an interface’s speed and duplex.

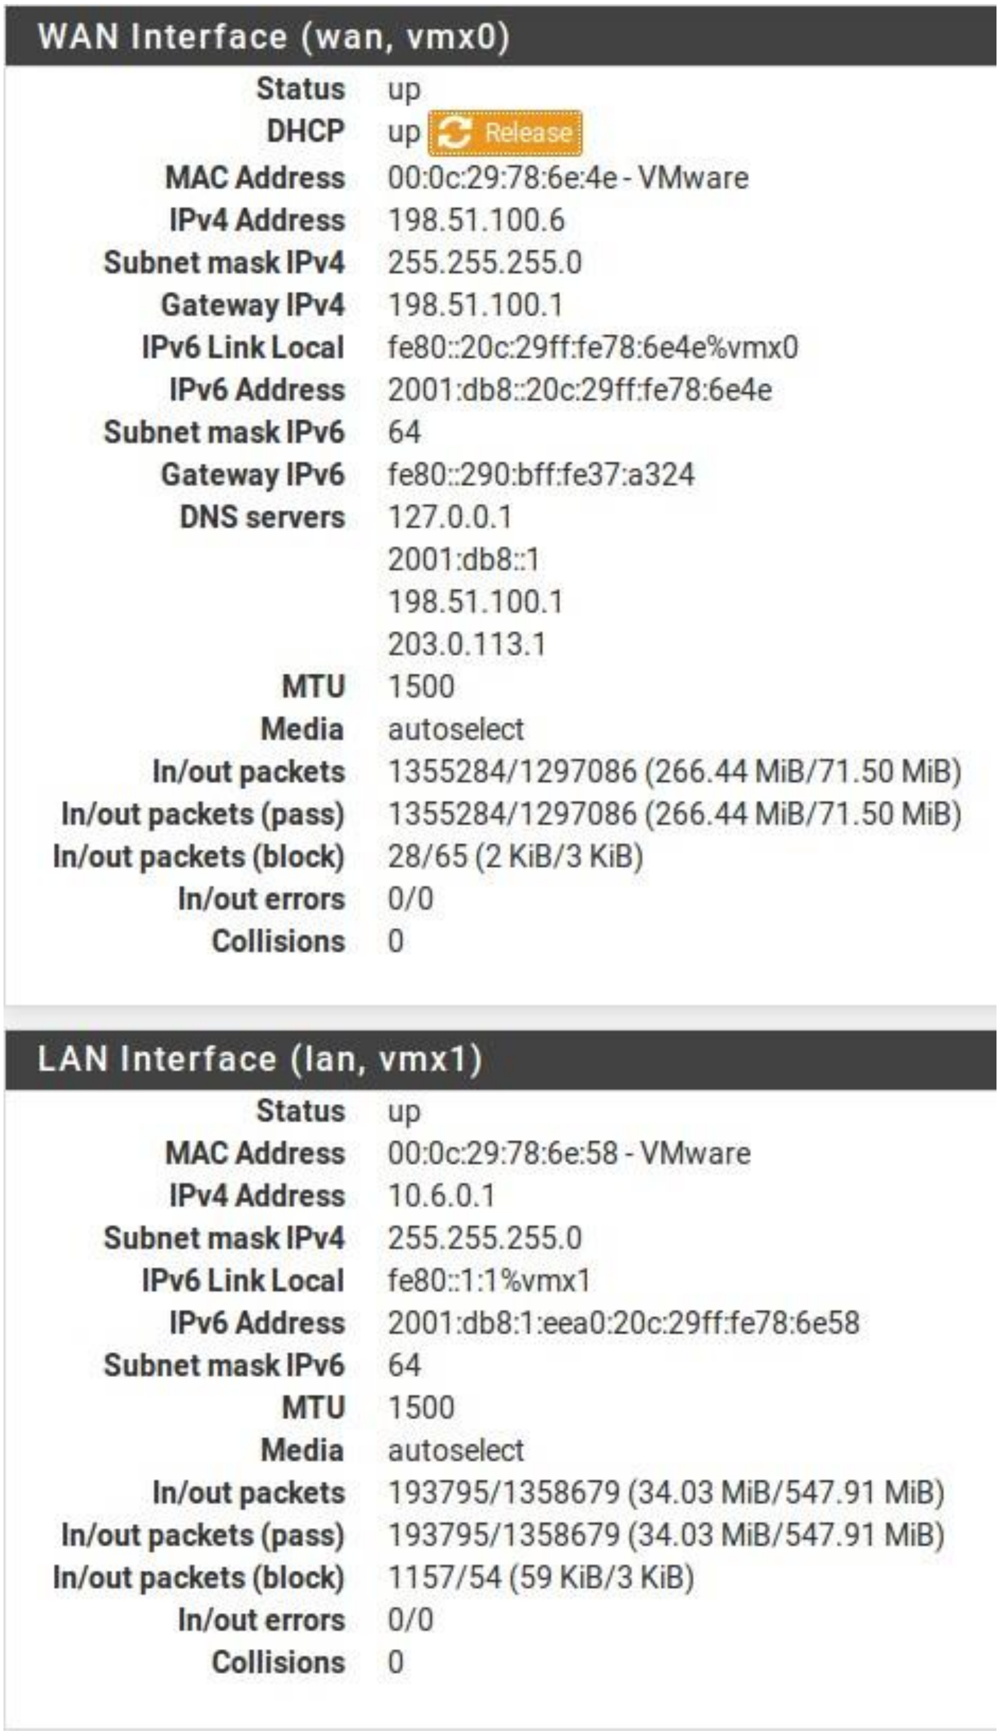

WAN Interface (wan, vmx0)

Status	up
DHCP	up Release
MAC Address	00:0c:29:78:6e:4e - VMware
IPv4 Address	198.51.100.6
Subnet mask IPv4	255.255.255.0
Gateway IPv4	198.51.100.1
IPv6 Link Local	fe80::20c:29ff:fe78:6e4e%vmx0
IPv6 Address	2001:db8::20c:29ff:fe78:6e4e
Subnet mask IPv6	64
Gateway IPv6	fe80::290:bff:fe37:a324
DNS servers	127.0.0.1 2001:db8::1 198.51.100.1 203.0.113.1
MTU	1500
Media	autoselect
In/out packets	1355284/1297086 (266.44 MiB/71.50 MiB)
In/out packets (pass)	1355284/1297086 (266.44 MiB/71.50 MiB)
In/out packets (block)	28/65 (2 KiB/3 KiB)
In/out errors	0/0
Collisions	0

LAN Interface (lan, vmx1)

Status	up
MAC Address	00:0c:29:78:6e:58 - VMware
IPv4 Address	10.6.0.1
Subnet mask IPv4	255.255.255.0
IPv6 Link Local	fe80::1:1%vmx1
IPv6 Address	2001:db8:1:eea0:20c:29ff:fe78:6e58
Subnet mask IPv6	64
MTU	1500
Media	autoselect
In/out packets	193795/1358679 (34.03 MiB/547.91 MiB)
In/out packets (pass)	193795/1358679 (34.03 MiB/547.91 MiB)
In/out packets (block)	1157/54 (59 KiB/3 KiB)
In/out errors	0/0
Collisions	0

Fig. 4: Interface Status

22.1.3 Service Status

Many system and package services show the status of their daemons at Status > Services.

Each service is shown with a name, a description, and the status, as seen in Figure *Services Status*. The status is listed as Running or Stopped.

Normally, it is not necessary to control services in this manner, but occasionally there are maintenance or troubleshooting reasons for doing so.

From this view, services can be controlled in various ways:

- Click ↻ to restart a running service

> Note: Some services will stop and start, others reload the configuration. Check the documentation of each service for details.

- Click ⏹ to stop a running service
- Click ↻ to start a stopped service

If available, other shortcuts are shown which navigate to pages related to the service. See *Quickly Navigate the GUI with Shortcuts* for information about shortcut icons.

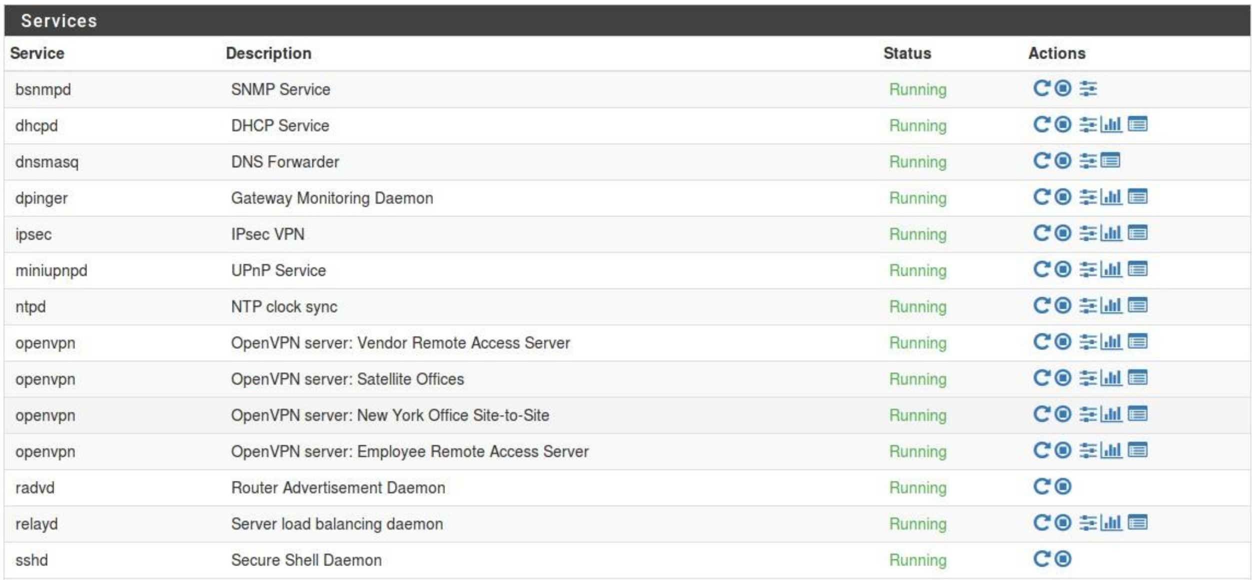

Services

Service	Description	Status	Actions
bsnmpd	SNMP Service	Running	↻ ⏹ ☰
dhcpd	DHCP Service	Running	↻ ⏹ ☰ 📊 ▤
dnsmasq	DNS Forwarder	Running	↻ ⏹ ☰ ▤
dpinger	Gateway Monitoring Daemon	Running	↻ ⏹ ☰ 📊 ▤
ipsec	IPsec VPN	Running	↻ ⏹ ☰ 📊 ▤
miniupnpd	UPnP Service	Running	↻ ⏹ ☰ 📊 ▤
ntpd	NTP clock sync	Running	↻ ⏹ ☰ 📊 ▤
openvpn	OpenVPN server: Vendor Remote Access Server	Running	↻ ⏹ ☰ 📊 ▤
openvpn	OpenVPN server: Satellite Offices	Running	↻ ⏹ ☰ 📊 ▤
openvpn	OpenVPN server: New York Office Site-to-Site	Running	↻ ⏹ ☰ 📊 ▤
openvpn	OpenVPN server: Employee Remote Access Server	Running	↻ ⏹ ☰ 📊 ▤
radvd	Router Advertisement Daemon	Running	↻ ⏹
relayd	Server load balancing daemon	Running	↻ ⏹ ☰ 📊 ▤
sshd	Secure Shell Daemon	Running	↻ ⏹

Fig. 5: Services Status

22.1.4 System Activity (Top)

The Diagnostics > System Activity page displays list of the top active processes running on the firewall. This is equivalent to running the command top -aSH at a shell prompt, except the GUI version does not have the CPU usage summary.

Using this view, it is easy to see processes that consume the most CPU power during a time of high load. For example, if the highest entry is an interrupt processing queue for one of the network cards, and the system isn't pushing enough traffic, it could be one sign that the firewall is trying to push more than the hardware can handle in the current configuration. If the top process is a PHP process, it could be that a browser has requested a GUI page that is processing a large amount of data.

Note: Threads that show idle in the COMMAND column indicate CPU time that is *not* in use (idle). It is normal for these to show 100% if the firewall has little to no load.

22.1.5 pfInfo

The Diagnostics > pfInfo page displays statistics and counters for the firewall packet filter which serve as metrics to judge how it is behaving and processing data. The information shown on the page contains items such as:

Bytes In/Out Bytes transferred in and out of the firewall.

Packets In/Out Packets transferred in or out and passed or blocked counters for each direction.

State Table / Source Tracking Table Statistics about the state table and source tracking table (*Firewall States*).

> Current Entries The number of entries in the table
>
> Searches How many times the table has been searched and the current rate of searches, which roughly corresponds to the number of packets being passed by the firewall on current open connections.
>
> Inserts The number of new states added to the table, and the rate at which the states are added. A high rate indicates that there are a lot of new connections being made to or through the firewall.
>
> Removals The number of old states being removed from the firewall.

Counters Statistics an counts for various types of special, unusual or badly formatted packets.

Limit Counters Counters that pertain to packets that have reached or exceeded limits configured on firewall rules, such as max states per IP address.

Table Size Limits State table max size, source node table size, frag table size, number of allowed tables, and maximum number of table entries.

State Timers The current configured timeout values for various connection states for TCP, UDP, and other protocols.

Interface Statistics Per-interface packet counters.

22.1.7 Filter Reload Status

The current status of a filter reload may be viewed in the AZTCO-FW® webGUI at Status > Filter Reload. A link to this page is available any time a filter change is made. The progress of the reload is displayed automatically, and it is updated automatically.

Normally, updates happen fast enough that Done. is the only message shown, indicating that there are no pending changes. With larger configurations, some delays are possible.

Press Reload Filter to initiate another filter reload. The progress will be updated as the process progresses.

When configuration synchronization is enabled, a Force Config Sync button is also present on this page. That button will trigger a forced synchronization of the firewall configuration by XMLRPC to the node specified by the *current configuration*.

22.1.8 Viewing the Contents of Tables

Aliases and other similar list of addresses are stored in a pf structure called a Table. These tables can be relatively static, as with the bogons list or aliases, or dynamic for things like snort or IP addresses exceeding connection limits. An alias becomes a "Table" once it has been loaded into the firewall ruleset. Tables may contain both IPv4 and IPv6 addresses, and the appropriate addresses are used based on the rules in which the tables are referenced.

The contents of these tables can be viewed at Diagnostics > Tables, which displays system and user-defined tables. On that page, select the desired table from the Table drop-down and the firewall will display its contents. If any alias contains a hostname, the contents of the alias are populated from DNS. Viewing the resulting table here confirms which IP addresses are in the table at that moment.

Individual entries may be removed by clicking 🗑 at the end of their row. Tables which are defined manually or by a file will be refreshed when the system performs a filter reload, so it is best to edit an alias and remove an entry rather than removing it from this page. Removing entries is best used for dynamic tables to remove an entry before it automatically expires.

Default Tables

The firewall includes several tables by default, depending on which features are enabled:

bogons/bogonsv6 If any interface is configured with *Block Bogon Networks* active, these tables will be present on the firewall. An 🔄 Update button is also presented for the bogon tables that will immediately re-fetch the bogons data rather than waiting for the usual monthly update.

tonatsubnets When using automatic outbound NAT, this table shows the list of networks for which automatic outbound NAT is being performed. Inspecting the table can aid in diagnosing tricky NAT issues to confirm if a subnet will have automatic outbound NAT applied to its traffic.

snort2c A dynamic table containing blocked offenders from IDS/IPS packages, Snort and Suricata.

virusprot A dynamic table containing addresses that have exceeded defined limits on firewall rules.

webConfiguratorlockout A dynamic table containing clients that repeatedly failed GUI login attempts.

sshlockout Similar to webConfiguratorlockout but used for tracking clients that fail repeated SSH login attempts.

22.1.9 Firewall States

AZTCO-FW® is a *stateful firewall* and uses one state to track each connection to and from the firewall. These states may be viewed in several ways in the WebGUI and from the console.

Viewing in the WebGUI

A listing of the firewall state table contents is available in the WebGUI by navigating to Diagnostics > States. Figure *Example States* shows a sample of the output displayed by the GUI.

The firewall displays several columns on this page, each with important information:

Interface The interface to which the state is bound. This is the interface through which the packet initially entered or exited the firewall.

Protocol The protocol of the traffic that created the state, such as TCP, UDP, ICMP, or ESP.

Source and Destination This column is in two parts, first the source, then an arrow indicating direction, and then the destination. The source and destination may also have a port number listed if the protocol in question uses ports. In cases where NAT is applied (outbound NAT, port forwards, or 1:1 NAT), the address is shown both before and after NAT has been applied.

For NAT such as outbound NAT which translates the source, the source section displays the translated source, and the original source inside parenthesis. For NAT types that translate the destination, such as port forwards, the destination section shows the translated destination and the original destination in parenthesis.

State The current status of the connection being tracked by this state entry. The specific values vary depending on the protocol. For example, TCP has many more state types than UDP or other connectionless protocols. The entry in this column contains two parts separated by a colon. The first part is the state for the source side, and the second part is the state for the destination side. See *Interpreting States* for more detail.

Packets The number of packets observed matching the state from the source and destination sides.

Bytes The total size of packets observed matching the state from the source and destination sides.

Individual states may be removed by clicking 🗑 at the end of their row.

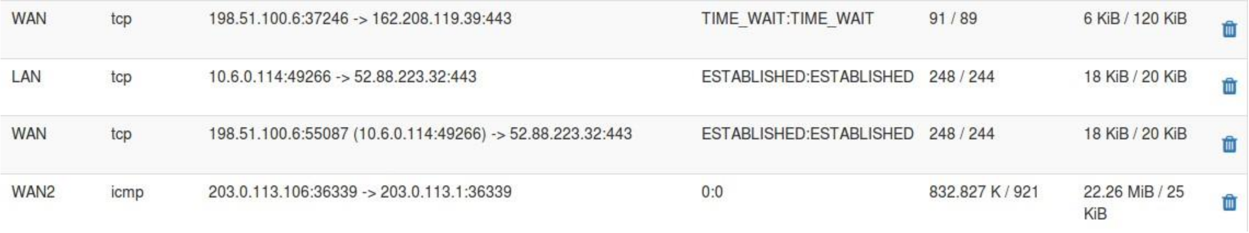

WAN	tcp	198.51.100.6:37246 -> 162.208.119.39:443	TIME_WAIT:TIME_WAIT	91 / 89	6 KiB / 120 KiB	🗑
LAN	tcp	10.6.0.114:49266 -> 52.88.223.32:443	ESTABLISHED:ESTABLISHED	248 / 244	18 KiB / 20 KiB	🗑
WAN	tcp	198.51.100.6:55087 (10.6.0.114:49266) -> 52.88.223.32:443	ESTABLISHED:ESTABLISHED	248 / 244	18 KiB / 20 KiB	🗑
WAN2	icmp	203.0.113.106:36339 -> 203.0.113.1:36339	0:0	832.827 K / 921	22.26 MiB / 25 KiB	🗑

Fig. 6: Example States

See also:

Firewall Maximum States

Filtering States

The State Filter panel enables quick searching of the state table contents to find items of interest.

To search for a state:

- Select a specific Interface in the State Filter panel or leave it on *all* to match all interfaces.
- Enter a Filter Expression which is a simple string of text to match exactly in the entry. Regular expressions are not supported in this field.
- Click Filter to locate the results.

All columns are searched for matching text, and only entries matching the text are displayed.

Tip: Searching for an IP address or subnet will also present a Kill States button which, when clicked, will remove all states originating from or going to the entered IP address or subnet.

Interpreting States

The State column for each state table entry provides information necessary to determine exactly what is happening with the connection. Each state entry contains two values with a colon between them, marking which value represents the state of the source (left), and which represents the destination (right).

A few of the most common state types are:

SYN_SENT For TCP connections, this indicates that the side showing this state sent a TCP SYN packet attempting to start a connection handshake.

CLOSED For TCP connections, the side with this status considers the connection closed, or no traffic has been received.

ESTABLISHED A TCP connection is considered fully established by this side.

TIME_WAIT/FIN_WAIT A TCP connection is in the process of closing and finishing up.

NO_TRAFFIC No packets have been received that match the state from this side.

SINGLE A single packet has been observed on this state from this side.

MULTIPLE Multiple packets have been observed on this state from this side.

Common pairings frequently found in the state table include:

ESTABLISHED:ESTABLISHED A fully established two-way TCP connection.

SYN_SENT:CLOSED The side showing *SYN_SENT* has sent a TCP SYN packet but no response has been received from the far side. Often this is due to the packet not reaching its destination, or being blocked along the way.

SINGLE:NO_TRAFFIC Similar to the above, but for UDP and other connectionless protocols. No response has been received from the destination side.

SINGLE:MULTIPLE For UDP and other connectionless protocols, commonly observed with DNS where the client sends one packet but receives a large response in multiple packets.

MULTIPLE:MULTIPLE For UDP and other connectionless protocols, there are multiple packets in both directions, which is normal for a fully operational UDP connection.

0:0 Indicates that there is no state level data. Typically only found on ICMP states, since ICMP does not have state levels like other protocols.

States Summary

The State Table Summary, accessible from Diagnostics > States Summary, provides statistics generated by an in-depth analysis of the state table and the connections therein.

The report includes the IP address, a total state count, and breakdowns by protocol and source/destination ports. Hovering over the ports shows a tooltip display of the full port list instead of the total number of ports. Depending on the firewall environment, high values by any metric may be normal.

The report includes the following categories:

By Source IP Address States summarized by the source IP address. This is useful for finding a potential source of attack, or a port scan or similar type probe/attack.

By Destination IP Address States summarized by the destination IP address of the connection. Useful for finding the target of an attack or identifying servers.

Total per IP Address States summarized by all connections to or from an IP address. Useful for finding active hosts using lots of ports, such as bittorrent clients.

By IP Address Pair Summarizes states between two IP addresses involved in active connections. Useful for finding specific client/server pairs that have unusually high numbers of connections.

> Warning: The States Summary can take a long time to process and display, especially if the firewall has an exceptionally large state table or a slow processor. In cases where the state table is extremely large, the page may not display properly or the page may fail with a memory error. In these cases, the summary page cannot be used.

Source Tracking States

When using Sticky Connections, the firewall maintains a source tracking table that records mappings of internal IP addresses to specific external gateways for connections that were passed by a rule utilizing a Load Balancing gateway group (Multiple gateways on the same tier). By default these associations only exist so long as there are active states from the internal IP address. There is a configurable timeout for these source tracking entries to allow them to exist longer if necessary.

See also:

For additional information about Sticky Connections and their related options, see config-advanced-sticky.

The source tracking associations are shown on Diagnostics > States on the Source Tracking tab, which is only visible if Sticky Connections are enabled.

The Source Tracking page lists the following information:

Source-to-Destination The mapping of a local IP address to a specific load balanced gateway.

# States The number of states matching this source IP address to any destination, including traffic that is not load balanced.

# Connections The number of states matching this source IP address which utilize the gateway. For example, connections leaving from this source to an Internet host.

Rate The rate of packets matching this source tracking entry.

These associations can be individually removed by clicking the Remove button at the end of each row.

Reset State Table / Source Tracking Table

Certain situations call for resetting the state table to force all existing connections to close and reestablish. The most notable examples are making changes to NAT rules, firewall block rules, or traffic shaping. When these types of changes are made, resetting the state table is the only way to make sure all connections respect the new ruleset or traffic shaping queues.

> Warning: Resetting the state table is disruptive, but clients may immediately reconnect provided they are still passed by the current firewall rules.

Both the state table and the source tracking table may be reset from Diagnostics > States on the Reset States tab. To

reset the tables, check either State Table, Source Tracking, or both, and then click 🗑 Reset.

> Warning: The browser will appear to lose connection with the firewall when resetting the state table. Once the browser realizes the old connection is invalid, it will reconnect. Close and reopen the browser to reconnect faster.

Viewing States with pfTop

pfTop is available from the GUI and the system console menu, and offers live views of the firewall ruleset, state table information, and related statistics.

pfTop in the GUI

In the GUI, pfTop can be found at Diagnostics > pfTop. The GUI offers several options to control the output:

View Controls the type of output displayed by pfTop. Not all views will contain meaningful information for every firewall configuration.

Default Shows a balanced amount of information, based around the source and destination of the traffic.

Label Centered around firewall rule descriptions.

Long Similar to the default view, but tailored for wider displays with longer rows for more columns of information. Shows the gateway after the destination.

Queue Shows the ALTQ traffic shaping queues and their usage.

Rules Shows firewall rules and their usage.

Size Shows states that have passed the most data.

Speed Shows states that have high-rate traffic.

State Shows status of states.

Time Shows long-lived states.

Sort By Some views can be sorted. When sorting is possible, the following sort methods are available. When selected, the view is sorted by the chosen column in descending order:

None No sorting, the natural order shown by the chosen view.

Age The age of the states.

Bytes The amount of data sent matching states.

Destination Address The destination IP address of the state.

Destination Port The destination port number of the state.

Expiry The expiration time of the state. This is the countdown timer until the state will be removed if no more data matches the state.

Peak The peak rate of traffic matching a state in packets per second.

Packet The number of packets transferred matching a state.

Rate The current rate of traffic matching a state in packets per second.

Size The total amount of traffic that has matched a state.

Source Port The source port number of the state.

Source Address The source IP address of the state.

Maximum # of States On views that support sorting, this option limits the number of state entries shown on the page.

pfTop on the Console

To access pfTop from the console or via ssh, use option 9 from the menu or run pftop from a shell prompt.

While viewing pfTop in this way, there are several methods to alter the view while watching its output. Press h to see a help screen that explains the available choices. The most common uses are using 0 through 8 to select different views, space for an immediate update, and q to quit. See the previous section for details on the meaning of the available views and sort orders.

The output is dynamically sized to the terminal width, with wider terminals showing much more information in additional columns.

22.1.10 Status

The status of the DHCP server service itself is at Status > Services. If the DHCP server is enabled, its status will be shown as Running, as in Figure *DHCP Daemon Service Status*. The buttons on the right side allow restarting or stopping the DHCP server daemon. Restarting is not normally necessary as AZTCO-FW® will automatically restart the service when configuration changes are made that require a restart. Stopping the service is also likely not necessary, as the service will stop when all instances of the DHCP server are disabled.

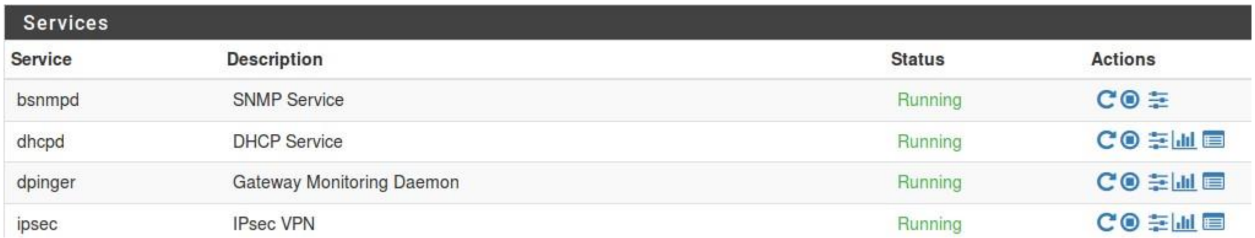

Services

Service	Description	Status	Actions
bsnmpd	SNMP Service	Running	↻ ◉ ☰
dhcpd	DHCP Service	Running	↻ ◉ ☰ 📊 ▤
dpinger	Gateway Monitoring Daemon	Running	↻ ◉ ☰ 📊 ▤
ipsec	IPsec VPN	Running	↻ ◉ ☰ 📊 ▤

Fig. 7: DHCP Daemon Service Status

22.1.11 Leases

The currently assigned DHCP leases are viewable at Status > DHCP leases. This page shows various aspects of the client leases. These include:

- The assigned IP address
- The client MAC address
- The hostname (if any) that the client sent as part of the DHCP request
- The description for a host with a DHCP static mapping
- The beginning and end times of the lease
- Whether or not the machine is currently online (in the ARP table)
- Whether or not the lease is active, expired, or a static registration

View inactive leases

By default, only active and static leases are shown, but everything, including the expired leases, may be displayed by clicking Show all configured leases. To reduce the view back to normal, click Show active and static leases only.

Wake on LAN Integration

Clicking the ⏻ icon to the right of the lease sends a Wake on LAN (WOL) packet to that host. Click ⊞ to create a WOL entry for the MAC address instead. For more details about Wake on LAN, see *Wake on LAN*.

Add static mapping

To create a static mapping from a dynamic lease, click ⊞ the to the right of the lease. This will pre-fill the MAC address of that host into the Edit static mapping screen. Add the desired IP address, hostname and description and click Save.

Delete a lease

While viewing the leases, an inactive or expired lease may be manually deleted by clicking 🗑 at the end of its line. This option is not available for active or static leases, only for offline or expired leases.

22.1.12 DHCP Service Logs

The DHCP daemon will log its activity to Status > System Logs, on the DHCP tab. Each DHCP request and response will be displayed, along with other status and error messages.

22.1.13 Viewing DHCPv6 Leases

A list of active and inactive DHCPv6 Leases (DHCP leases for IPv6 hosts) and delegated prefixes can be viewed in AZTCO-FW® software by navigating to Status > DHCPv6 Leases.

All active leases are shown, along with the IPv6 address, IAID, DUID, MAC address, hostname, lease start and end times, lease type, and whether or not the system is online. (As with the *NDP Table*, this is not always a reliable indicator)

IPv6 hosts are identified by a combination of the Interface Association Identifier (IAID) and the DHCP Unique Identifier (DUID). The DUID is meant to be unique *per device* and the IAID is unique *per interface on the device*. Due to variances in DHCPv6 clients between operating systems and manufacturers some clients do not send an IAID, and some DUID values are sized differently than others.

To view expired leases, click Show All Configured Leases. To switch the view back, click Show Active and Static Leases.

A DHCPv6 static IP mapping may be added by clicking ✚. An offline dynamic lease may be deleted by clicking 🗑.

Delegated Prefixes

When Prefix Delegation is enabled, the bottom of the page lists delegated prefixes and their routing. Similar to the lease status, the DUID of the target system is shown along with the start and end time of the delegation. The IPv6 Prefix column shows the delegated prefixes, their prefix length, and the address to which they are routed. The target address and DUID will match one of the leases in the list at the top of the screen.

22.1.14 Gateway Status

In the AZTCO-FW® webGUI, Status > Gateways displays the current status of all configured gateways.

The status output includes the gateway name, gateway IP, Monitor IP, status and description. The status field will show Online, Offline, or Warning, and a textual description of the problem (or success).

The Gateway Groups tab shows the status of gateway group members and the groups as a whole.

22.1.15 CARP Status

The CARP status page located through the AZTCO-FW® webGUI at Status > CARP (failover) shows the current status of all configured CARP *Virtual IP addresses*. It also provides some controls to enable and disable CARP for troubleshooting and maintenance.

For each VIP, the Interface, Virtual IP, and Status are shown.

The Interface column shows the interface name or identifier for the VIP at the operating system level. This has changed over the years, and currently shows as "_vip" (ex: *lan_vip1*) on AZTCO-FW software version 2.1.x, and on AZTCO-FW software version 2.2, it shows as "@" (ex: *LAN@1*).

The Virtual IP column shows the IP address associated with the entry.

The Status column shows MASTER, BACKUP, or INIT, which are:

- MASTER: Indicates this node is accepting all traffic for this VIP
- BACKUP: Indicates this node is monitoring CARP advertisements and not accepting traffic for the VIP.
- INIT: Generally indicates a problem with the VIP. Either the VIP is not configured at the OS level, or the interface upon which it is configured is down or has a problem.

A toggle button at the top of the list will show either Disable CARP or Enable CARP depending on the current status. Disabling CARP will remove the VIPs from the system, and the next available node will take over as MASTER for the CARP VIPs. This setting is not remembered across reboots.

On the primary node, each VIP should show MASTER. On the secondary node, each VIP should show BACKUP. If both nodes show MASTER, there is usually a problem at layer 2 (the switch) preventing the two nodes from seeing each others' advertisements. See *CARP Configuration Troubleshooting*.

Maintenance Mode

There is a toggle button to Enter Persistent CARP Maintenance Mode or Leave Persistent CARP Maintenance Mode. This mode persists across reboots, so it is useful for performing maintenance or upgrades on the primary node such that it does not cause it to take back over prematurely before it is ready.

Widget

This is also a CARP Status widget available for the *Dashboard* that shows the same information in a condensed format without the control buttons.

22.1.16 Captive Portal Status

The Captive Portal Status page is available through the AZTCO-FW® software webGUI at Status > Captive Portal.

To view the status for a Zone, select it from the Captive Portal Zone drop-down.

This page displays a list of online users, including their IP address, MAC address, Username, and Session Start Time.

The exact output for each field depends on the portal configuration. For example, if MAC address filtering is disabled, the MAC address may not be displayed. If user authentication is not required, then there will be no Username to display.

One of several different user styles may appear in the list.

For a portal with No Authentication, a listing like Figure *Online Captive Portal Users No Authentication* is shown.

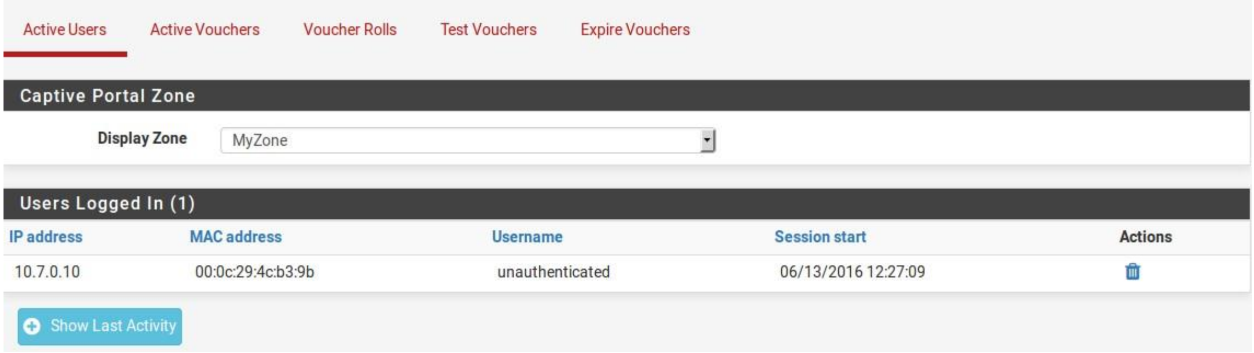

Captive Portal Zone

Display Zone: MyZone

Users Logged In (1)

IP address	MAC address	Username	Session start	Actions
10.7.0.10	00:0c:29:4c:b3:9b	unauthenticated	06/13/2016 12:27:09	🗑

Fig. 8: Online Captive Portal Users No Authentication

For Local User Manager or RADIUS authentication is selected, a listing like Figure *Online Captive Portal Users User Authentication* is shown. If RADIUS with MAC Authentication is used, the username will be the MAC address.

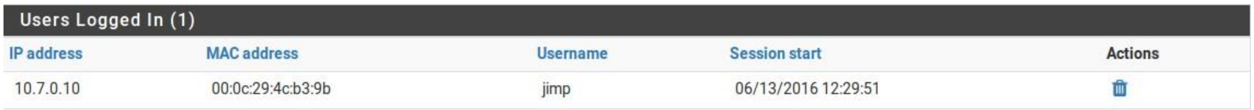

Users Logged In (1)

IP address	MAC address	Username	Session start	Actions
10.7.0.10	00:0c:29:4c:b3:9b	jimp	06/13/2016 12:29:51	🗑

Fig. 9: Online Captive Portal Users User Authentication

If vouchers are active, users that signed on using a voucher will show like Figure *Online Captive Portal Users Vouchers*.

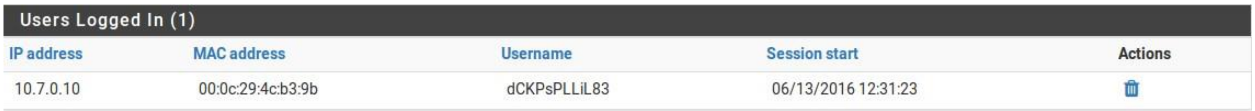

Users Logged In (1)

IP address	MAC address	Username	Session start	Actions
10.7.0.10	00:0c:29:4c:b3:9b	dCKPsPLLiL83	06/13/2016 12:31:23	🗑

Fig. 10: Online Captive Portal Users Vouchers

22.1.17 Viewing Active Network Sockets

The Diagnostics > Sockets page presents a list of active TCP/IP sockets for both IPv4 and IPv6 *used by the firewall itself*.

This list is useful for determining which IP addresses and ports are in use by various system processes or packages. It is interpreted from the output of the FreeBSD command "*sockstat*".

By default, only *listening* sockets are shown. Click Show all socket connections to also display sockets in use by the system making connections to external hosts.

The output of this command only shows sockets used by the firewall OS for daemons or other programs on the firewall. It does not show connections for traffic passing through the firewall.

22.1.18 ARP Table

ARP (Address Resolution Protocol) is used for locating IPv4 systems on a local network by MAC address.

The ARP table in AZTCO-FW® software displays a list of systems on the network that have attempted to talk to or through the AZTCO-FW firewall within the past few minutes. If a system is up but has not talked to (or through) the AZTCO-FW firewall it will not show up in the ARP table.

To view the list of systems currently seen by AZTCO-FW software, click Diagnostics > ARP Table. This list shows the IP Address, MAC Address, Hostname and the Interface where each system was last seen.

For IPv6 hosts, see *NDP Table*.

22.1.19 NDP Table

The NDP Table, found under the Diagnostics menu in the AZTCO-FW® webGUI, shows the IPv6 Neighbor Discovery Protocol list. This list contains all of the current IPv6 peers, and is roughly analogous to the *ARP Table* for IPv4.

The NDP table shows the peer's IPv6 address, its MAC address, the Hostname (if known), and the interface upon which that peer IP address was last observed.

As with ARP, these are only the hosts that are actively talking to or through the firewall, and is not a perfect indicator of a host's actual online status. If a host is in the table, it is likely online, but if it is not in the list, it does not necessarily mean that it is offline, it may simply not have attempted external communication recently.

22.1.20 IPsec Status

The status of all IPsec tunnels may be viewed by browsing to Status > IPsec. This page is divided into four tabs.

Overview Tab

This tab lists all enabled IPsec tunnels, the local and remote IP addresses, local and remote networks, tunnel description, and status.

A green icon indicates that the tunnel is up (has SAD and SPD entries, signifying a complete phase 1 and 2 connection).

A yellow icon indicates that the tunnel is not fully up and active.

SAD Tab

Shows the contents of the IPsec Security Association Database. There should be one for each “direction” between *public peer addresses* of an active IPsec tunnel.

SPD Tab

Shows the contents of the IPsec Security Policy Database. There should be one for each direction between *private networks* of an active IPsec tunnel.

Logs Tab

A shortcut to the *IPsec Logs* normally found at Status > System Logs, IPsec tab.

See Also

- *IPsec Troubleshooting*

22.1.21 Checking the Status of OpenVPN Clients and Servers

The OpenVPN status page at Status > OpenVPN shows the status of each OpenVPN server and client. Service start/stop controls are also available for each separate server and client instance on the status page.

For OpenVPN servers in SSL/TLS server mode, the status provides a list of connected remote clients along with their usernames or certificate common names, as seen in Figure *OpenVPN Status for SSL/TLS Server With One Connected Client*. Clients may also be disconnected from this screen by clicking the ✖ at the end of the client row. For these servers a ⊕ Show Routing Table button is also displayed. Clicking this button will show a table of networks and IP addresses connected through each client certificate.

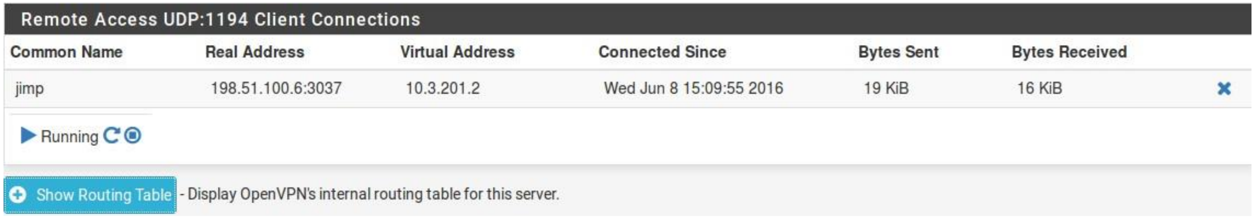

Remote Access UDP:1194 Client Connections

Common Name	Real Address	Virtual Address	Connected Since	Bytes Sent	Bytes Received	
jimp	198.51.100.6:3037	10.3.201.2	Wed Jun 8 15:09:55 2016	19 KiB	16 KiB	✖

▶ Running

⊕ Show Routing Table - Display OpenVPN's internal routing table for this server.

Fig. 11: OpenVPN Status for SSL/TLS Server With One Connected Client

For OpenVPN servers in shared key mode, the status will indicate whether it’s running and waiting on connections, or if the remote client has connected.

For OpenVPN clients, the status indicates whether a connection is pending or active.

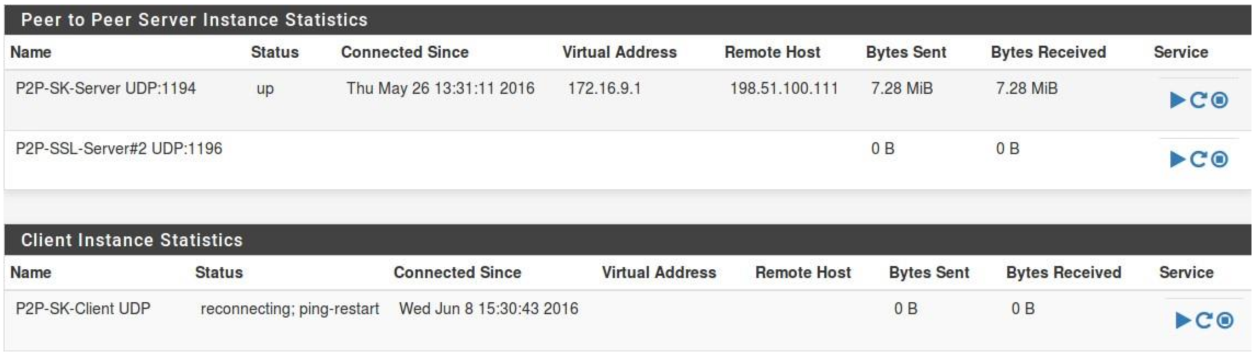

Peer to Peer Server Instance Statistics

Name	Status	Connected Since	Virtual Address	Remote Host	Bytes Sent	Bytes Received	Service
P2P-SK-Server UDP:1194	up	Thu May 26 13:31:11 2016	172.16.9.1	198.51.100.111	7.28 MiB	7.28 MiB	▶⟳◉
P2P-SSL-Server#2 UDP:1196					0 B	0 B	▶⟳◉

Client Instance Statistics

Name	Status	Connected Since	Virtual Address	Remote Host	Bytes Sent	Bytes Received	Service
P2P-SK-Client UDP	reconnecting; ping-restart	Wed Jun 8 15:30:43 2016			0 B	0 B	▶⟳◉

Fig. 12: OpenVPN Status Showing a Server that is up, one waiting for a Connection, and a Client Attempting to Reconnect

22.1.22 Route Table Contents

View the firewall route table at Diagnostics > Routes.

Each line shows the destination network, gateway by which it can be reached, route flags, number of references and times the route has been used, the MTU, the interface through which the routed traffic will be sent by the firewall.

The list of routes supports pagination and filtering to aid with viewing large routing tables such as those found with a full BGP feed.

Resolve Names When checked, the firewall attempts a DNS lookup to show hostnames rather than IP addresses for route table entries. This will cause a delay and performance penalty.

Rows to display By default the page displays 100 rows. Choose a new value to show more or less rows.

Filter Search the route table for entries matching this string. The field supports regular expressions for advanced filtering.

Click ⟳ Update to redisplay the routing table with the current settings.

The meaning of the Flags column entries are explained in *Route Table Flags*

Table 1: Route Table Flags

Letter	Flag	Meaning
1	RTF_PROTO1	Protocol specific routing flag #1
2	RTF_PROTO2	Protocol specific routing flag #2
3	RTF_PROTO3	Protocol specific routing flag #3
B	RTF_BLACKHOLE	Just discard pkts (during updates)
b	RTF_BROADCAST	The route represents a broadcast address
C	RTF_CLONING	Generate new routes on use
c	RTF_PRCLONING	Protocol-specified generate new routes on use
D	RTF_DYNAMIC	Created dynamically (by redirect)
G	RTF_GATEWAY	Destination requires forwarding by intermediary
H	RTF_HOST	Host entry (net otherwise)

L	RTF_LLINFO	Valid protocol to link address translation
M	RTF_MODIFIED	Modified dynamically (by redirect)
R	RTF_REJECT	Host or net unreachable
S	RTF_STATIC	Manually added
U	RTF_UP	Route usable
W	RTF_WASCLONED	Route was generated as a result of cloning
X	RTF_XRESOLVE	External daemon translates proto to link address

22.1.23 Monitoring the Queues

Monitor the shaper using Status > Queues to ensure that traffic shaping is working as intended. As can be seen in Figure *Basic WAN Queues*, this screen shows each queue listed by name, its current usage, and other related statistics.

Queue The name of the traffic shaper queue.

Statistics A graphical bar which shows how “full” this queue is.

PPS The rate of queued data in packets per second (PPS)

Bandwidth The rate of queued data in bits per second (e.g. Mbps, Kbps, bps).

Borrows Borrows happen when a neighboring queue is not full and capacity is borrowed from there.

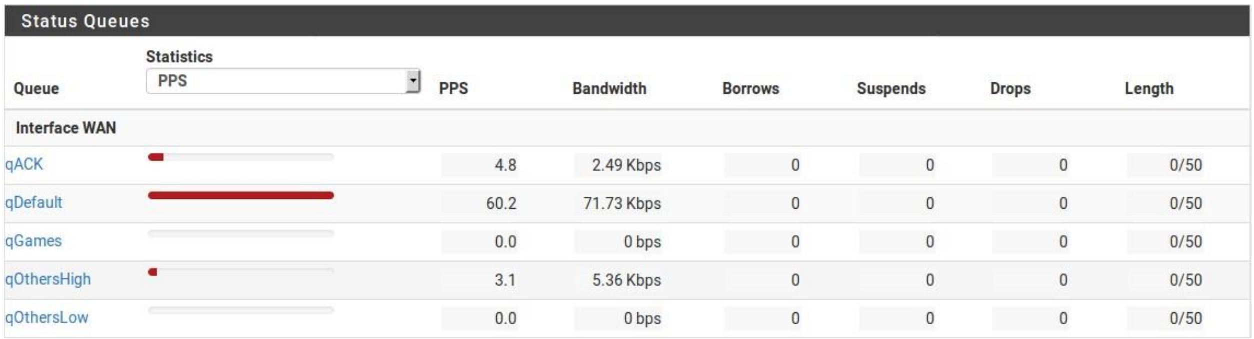

Status Queues

Queue	Statistics PPS	PPS	Bandwidth	Borrows	Suspends	Drops	Length
Interface WAN							
qACK		4.8	2.49 Kbps	0	0	0	0/50
qDefault		60.2	71.73 Kbps	0	0	0	0/50
qGames		0.0	0 bps	0	0	0	0/50
qOthersHigh		3.1	5.36 Kbps	0	0	0	0/50
qOthersLow		0.0	0 bps	0	0	0	0/50

Fig. 13: Basic WAN Queues

Suspends The suspends counter indicates when a delay action happens. The suspends counter is only used with the CBQ scheduler and should be zero when other schedulers are in use.

Drops Drops happen when traffic in a queue is dropped in favor of higher priority traffic. Drops are normal and this does not mean that a full connection is dropped, only a packet. Usually, one side of the connection will see that a packet was missed and then resend, often slowing down in the process to avoid future drops.

Length The number of packets in the queue waiting to be transmitted, over the total size of the queue.

22.1.24 Status

The NTP status page shows the status of each NTP peer server. This status page can be found at Status > NTP. An example of the status is shown in Figure *NTP Daemon Status With GPS Output*.

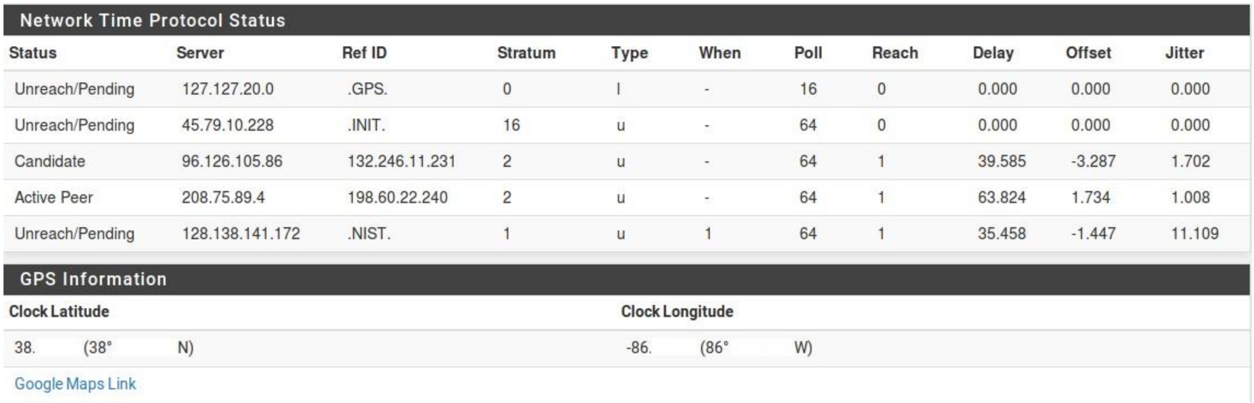

Network Time Protocol Status

Status	Server	Ref ID	Stratum	Type	When	Poll	Reach	Delay	Offset	Jitter
Unreach/Pending	127.127.20.0	.GPS.	0	l	-	16	0	0.000	0.000	0.000
Unreach/Pending	45.79.10.228	.INIT.	16	u	-	64	0	0.000	0.000	0.000
Candidate	96.126.105.86	132.246.11.231	2	u	-	64	1	39.585	-3.287	1.702
Active Peer	208.75.89.4	198.60.22.240	2	u	-	64	1	63.824	1.734	1.008
Unreach/Pending	128.138.141.172	.NIST.	1	u	1	64	1	35.458	-1.447	11.109

GPS Information

Clock Latitude	Clock Longitude
38. (38° N)	-86. (86° W)

Google Maps Link

Fig. 14: NTP Daemon Status With GPS Output

The status screen contains one line for every peer, and lists the peer IP address or server ID, the reference clock ID for the peer and various other values that indicate the general quality of the NTP server from the perspective of this firewall. The first column is the most useful, as it indicates which peer is currently the active peer for time sync, which servers are potential candidates to be peers, and which servers have been rejected and why.

If a serial GPS is connected and configured, the coordinates reported by the GPS device are also listed, along with a link to the coordinates on Google Maps.

22.2 Graphs

These articles cover graphs for monitoring AZTCO-FW itself as well as for traffic on interfaces and using additional packages for more detailed monitoring of user throughput/usage.

22.2.1 Monitoring Graphs

AZTCO-FW® software has many built-in graphs that monitor different aspects of the system, and they work out-of-the-box with no intervention.

The firewall collects and maintains data about how the system performs, and then stores this data in Round-Robin Database (RRD) files. Graphs created from this data are available under Status > Monitoring.

These graphs measure things such as CPU usage, memory usage, state table usage, throughput (in bytes as well as packets), link quality, traffic shaping queue usage, and more.

The graph on that page can be configured to show items from several categories, and a category and graph may be chosen for both the left axis and right axis for easy comparison.

Working with Graphs

The firewall displays a graph showing its CPU usage by default. To view other graphs or to add a second category on another axis, the graph settings must be changed as described in the next section, *Graph Settings*.

Inside the graph, the labels in the top left corner note the sources for the data in the left axis and right axis.

The graph contains a legend at the top right with each of the data sources plotted on the graph. Clicking a data source in the legend will hide it from view.

Tip: If a data source has a large spike, click its name in the legend to remove it from the graph. With the larger data source removed, more detail from the other remaining sources will be visible.

The firewall hostname, graph time period, and graph resolution are printed along the bottom of the graph, along with the time the graph was generated.

The firewall prints a table below the graph itself with a summarization of the data. This table contains minimums, averages, maximums, current values, in some cases 95th percentile values. In cases where units are given, hovering the mouse pointer over the unit will display a more detailed description of the unit.

Note: Totals are not displayed because the way data is stored in RRD files, accurate totals are not possible. To see total usage for traffic on network interfaces, install the Status Traffic Totals package.

Figure *WAN Traffic Graph* shows an example of an 8-hour graph of traffic on a firewall interface named *CABLE* with inverse enabled. The interface has a maximum utiliztion of 9.96Mbit/s during a 1 minute period.

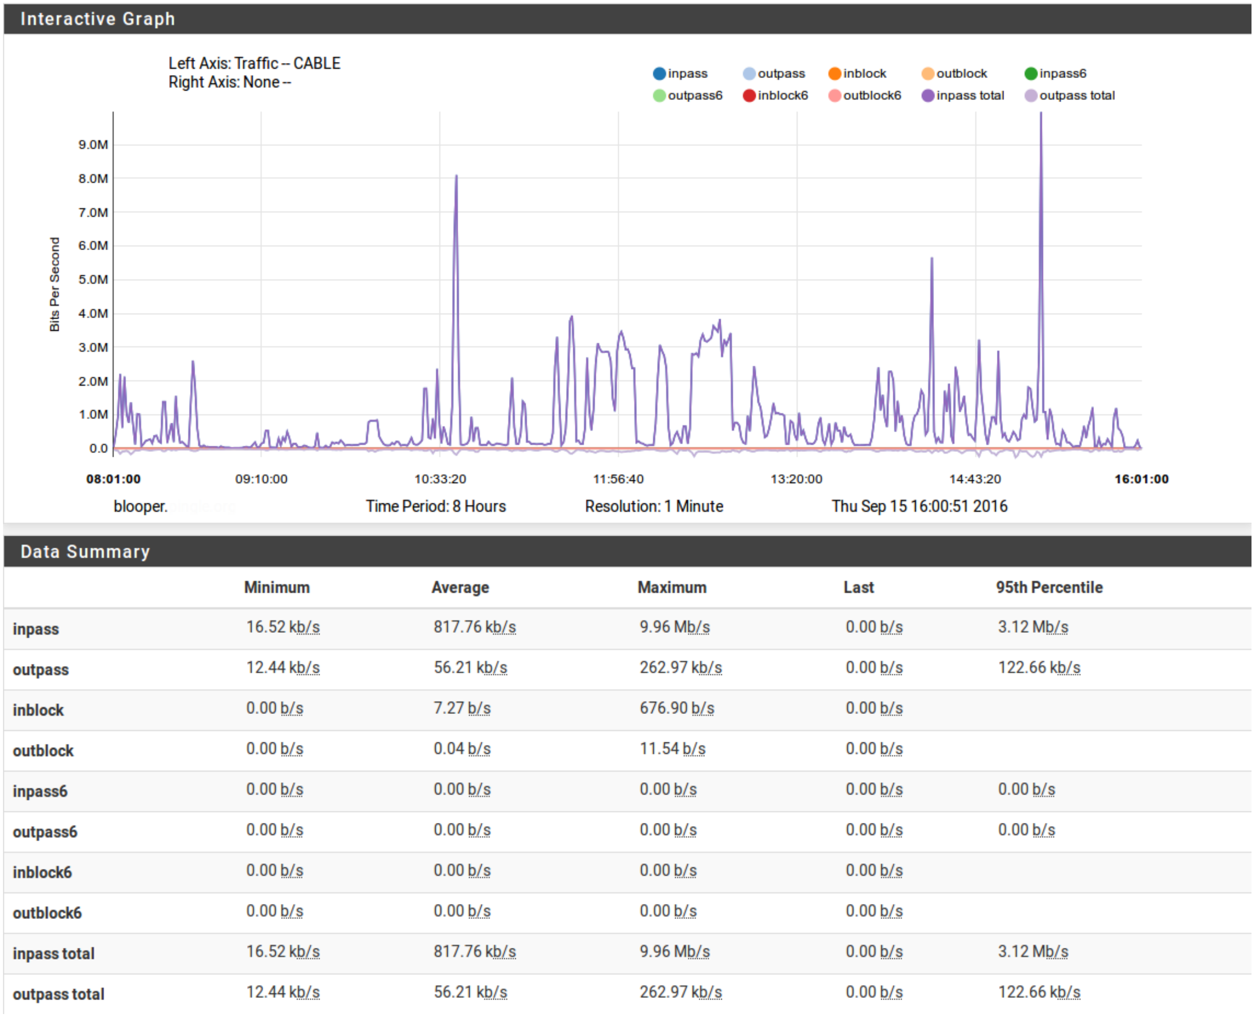

Left Axis: Traffic -- CABLE
Right Axis: None --

Time Period: 8 Hours

Resolution: 1 Minute

Thu Sep 15 16:00:51 2016

	Minimum	Average	Maximum	Last	95th Percentile
inpass	16.52 kb/s	817.76 kb/s	9.96 Mb/s	0.00 b/s	3.12 Mb/s
outpass	12.44 kb/s	56.21 kb/s	262.97 kb/s	0.00 b/s	122.66 kb/s
inblock	0.00 b/s	7.27 b/s	676.90 b/s	0.00 b/s	
outblock	0.00 b/s	0.04 b/s	11.54 b/s	0.00 b/s	
inpass6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s
outpass6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s
inblock6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	
outblock6	0.00 b/s	0.00 b/s	0.00 b/s	0.00 b/s	
inpass total	16.52 kb/s	817.76 kb/s	9.96 Mb/s	0.00 b/s	3.12 Mb/s
outpass total	12.44 kb/s	56.21 kb/s	262.97 kb/s	0.00 b/s	122.66 kb/s

Fig. 15: WAN Traffic Graph **Graph Settings**

To change the graph, click 🔧 in the breadcrumb bar to display the graph settings panel.

Tip: The graph settings panel is hidden by default but this behavior can be changed. Navigate to System > General Setup and check Monitoring Settings to always display the settings panel by default.

The options on the settings panel are:

Left Axis / Right Axis The options here control the data dispayed on each axis. By default only the Left Axis is populated with a value, but both may be utilized to compare areas. First pick a Category (or *None*), then pick a Graph inside that category. The list of available categories and graphs will vary depending on the firewall configuration.

Category The general area of the desired graph: System, Traffic, Packets, Quality, Captive Portal, NTP, Queues, QueueDrops, DHCP, Cellular, Wireless, and VPN Users. These are covered in more detail later in this section.

Graph The specific graph to display from the chosen category.

Options This section of the settings panel controls how the graph itself looks, including the time span and style.

Time Period The length of time to show on the graph. The default ranges cover from 1 hour up to 4 years, or a *Custom* period may be chosen. Selecting *Custom* displays the Custom Period controls. All of the periods are displayed even if there is no data in a graph database going back that far. The graph will be empty for times when graphing was not active.

Resolution The smallest slice of time for which data is available on this graph. Over time, data is consolodated over longer periods so resolution is lost. For example, on a 1-hour graph it is possible to see data from one minute intervals, but on a graph including older data, it is not possible to show data that accurately since it has been averaged out. Depending on the time period of the graph it may contain *1 Minute*, *5 Minute*, *1 Hour*, or *1 Day* averages for data. Resolutions which are not possible for the given time period cannot be selected.

Inverse Used on graphs such as the traffic graph, to separate incoming and outgoing data. For example, with Inverse set to *On*, outbound data is represented as a negative value to more easily differentiate it from inbound data.

Custom Period When Time Period is set to *Custom*, the GUI displays this section to configure the custom time period for the graph.

Start Date The start date for the graph. Clicking in the field will show a calendar date picking control. Only today, or days in the past, may be selected.

Start Hour The hour of the day to start the graph using 24-hour style (0-23).

End Date The end date for the graph.

End Hour The end hour for the graph, exclusive. The chosen hour is not included in the graph. For example, on a graph starting at hour 10 to hour 12, the graph covers 10:00am to 12:00pm.

Settings Click ⚙ Show Advanced to display additional advanced controls not typically required for average use.

Export as CSV Click this button to download the data from the graph as a .csv (Comma Separated Values) spreadsheet file, which can then be imported into another program for analysis.

Save as Defaults Click this button to store the current graph settings as the default configuration so this specific graph will be displayed by default on future visits to this page.

Disable/Enable Graphing This toggle will disable or enable the collection of graph data. Graphing is enabled by default. Normally this would only be disabled for diagnostic purposes or if all required graphing is handled externally.

Reset Graphing Data Clicking this button will erase all graph database files and create new, empty files.

Click 🔄 Update Graphs to change the graph to the selected view.

Graph Category List

There are a several different categories of graph data that the firewall can plot. Each category is covered here, but not all categories will be visible on every firewall. Some graphs must be enabled separately or will only be present if a specific feature or piece of hardware is enabled.

System Graphs

The graphs under the *System* category show a general overview of the system utilization, including CPU usage, memory usage, and firewall states.

Mbuf Clusters

The Mbuf Clusters graph plots the network memory buffer cluster usage of the firewall. Firewalls with many interfaces, or many CPU cores and NICs that use one interface queue per core, can consume a large number of network memory buffers. In most cases, this usage will be fairly flat, but depending on various circumstances, such as unusually high load, the values may increase. If the usage approaches the configured maximum, increase the number of buffers.

See also:

Refer to *Hardware Tuning and Troubleshooting* for information on how to increase the amount of mbufs available to the OS.

The Mbuf Clusters graph contains the following data sources:

Current The current number of consumed mbuf clusters

Cache The number of cached mbuf clusters

Total The total of Current and Cache

Max The maximum allowed number of mbuf clusters

Memory Graph

The Memory graph shows the system RAM usage broken down using the following data sources:

Active The amount of active (in use) memory

Inactive The amount of inactive memory, which was in use, but could be reallocated.

Free The amount of free memory, which is not used at all.

Cache The amount of memory used for caching by the operating system.

Wire The amount of wired memory, typically kernel memory

Note: The OS will attempt to use RAM as much as posssible for caching rather than allowing it to sit idle, so the amount of free RAM will often appear lower than expected. If memory demand increases, cached memory will be made available for use.

Processor Graph

The processor graph shows CPU usage for the firewall using the following data sources:

User Utilization The amount of processor time consumed by user processes.

Nice Utilization The amount of processor time consumed by processes with a high priority.

System Utilization The amount of processor time consumed by the operating system and kernel.

Interrupts The amount of processor time consumed by interrupt handling, which is processing hardware input and output, including network interfaces.

Processes The number of running processes.

States Graph

The states graph shows the number of system states but also breaks down the value in several ways.

State Changes The number of state changes per second, or “churn”. A high value from this source would indicate a rapid number of new or expiring connections.

Filter States The total number of state entries in the states table.

Source Addresses The number of active unique source IP addresses.

Destination Addresses The number of active unique destination IP addresses.

Traffic Graphs

Traffic graphs shows the amount of bandwidth used on each available interface in *bits per second* notation. The Graph list contains entries for each assigned interface, as well as IPsec and individual OpenVPN clients and servers.

The traffic graph is broken down into several data sources. Aside from the total, each has an IPv4 and IPv6 equivalent. The IPv6 data sources have 6 appended to the name.

inpass The rate of traffic entering this interface that was *passed* into the firewall.

outpass The rate of traffic leaving from this interface that was *passed* out of the

firewall. inblock The rate of traffic attempting to reach this interface that was *blocked*

from entering the firewall. outblock The rate of traffic attempting to leave this

interface that was *blocked* from leaving the fiewall. inpass total The total rate of traffic

(IPv4 and IPv6) that was passed inbound.

outpass total The total rate of traffic (IPv4 and IPv6) that was passed outbound.

Note: The terms “inbound” and “outbound” on these graphs are from the perspective of the firewall itself. On an external interface such as a WAN, “inbound” traffic is traffic arriving at the firewall from the Internet and “outbound” traffic is traffic leaving the firewall going to a destination on the Internet. For an internal interface, such as LAN, “inbound” traffic is traffic arriving at the firewall from a host on the LAN, likely destined for a location on the Internet and “outbound” traffic is traffic leaving the firewall going to a host on the LAN.

Packet Graphs

The packet graphs work much like the traffic graphs and have the same names for the data sources, except instead of reporting based on bandwidth used, it reports the number of *packets per second* (pps) passed. The Graph list contains entries for each assigned interface, as well as IPsec and individual OpenVPN clients and servers.

Packets Per Second (pps) is a better metric for judging hardware performance than Traffic throguhput as it more accurately reflects how well the hardware handles packets of any size. A circuit may be sold on a certain level of bandwidth, but hardware is more likely to be bottlenecked by an inability to handle a large volume of small packets. In situations where the hardware is the limiting factor, the Packets graph may show a high plateau or spikes while the traffic graph shows usage under the rated speed of the line.

Quality Graphs

The Quality category contains Graph entries that track the quality of WAN or WAN-like interfaces such as interfaces with a gateway specified or those using DHCP or PPPoE. The firewall contains one Graph entry per gateway, including gateways that were configured previously, but no longer exist. Graph data files for old gateways are not automatically removed so that historical data is available for future reference.

The following data sources are used to track gateway reliability:

Packet Loss The percentage of attempted pings to the monitor IP address that were lost. Loss on the graph indicates connectivity issues or times of excessive bandwidth use where pings were dropped.

Delay Average The average delay (Round-trip time, RTT) on pings sent to the monitor IP address. A high RTT means that traffic is taking a long time to make the round trip from the firewall to the monitor IP address and back. A high RTT could be from a problem on the circuit or from high utilization.

Delay Standard Deviation The standard deviation on the RTT values. The standard deviation gives an impression of the variability of the RTT during a given calculation period. A low standard deviation indicates that the connection is relatively stable. A high standard deviation means that the RTT is flucuating up and down over a large range of values, which could mean that the connection is unstable or very busy.

Captive Portal

The Captive Portal category contains Graph entries for each Captive Portal zone, past and present. Graph data files for old zones are not automatically removed.

Concurrent The *Concurrent* graph choice shows how many users are logged in at a given point in time. As users log out or their sessions expire, this count will go down. A large number of concurrent users will not necessarily cause a strain on the portal, but it can be useful for judging overall capacity and bandwidth needs.

Logged In The *Logged In* graph shows the number of login events that occur during each polling interval. This is useful for judging how busy the captive portal daemon is at a given point in time. A large number of users logging in around the same time will put more stress on the portal daemon compared to logins that are spread out over the course of a day.

NTP

The NTP graph displays statistics about the NTP service and clock quality. This graph is disabled by default because it is not relevant for most use cases. The graph can be enabled at Services > NTP. On that page, check Enable RRD Graphs of NTP statistics.

See also:

For more information about these values, see the NTP Configuration Manual, NTP Query Manual, and the NTPv4 Specification.

Offset Combined clock difference between from server relative to this host.

System Jitter (sjit) Combined system jitter, which is an estimate of the error in determining the offset.

Clock Jitter (cjit) Jitter computed by the clock discipline module.

Clock Wander (wander) Clock frequency stability expressed in parts per million (PPM)

Frequency Offset (freq) Offset relative to hardware clock (In PPM)

Root Dispersion (disp) Total difference between the local clock and the primary reference clock across the network.

Queue/Queuedrops Graphs

The queue graphs are a composite of each traffic shaper queue. Each individual queue is shown, represented by a unique color.

The Queues category shows individual queue usage in bytes.

The QueueDrops category shows a count of packet drops from each queue.

DHCP

The DHCP category contains a graph for each interface with a DHCP server enabled. The data sources shown for DHCP are:

Leases The number of leases in use out of the configured DHCP range for the interface.

Static Leases The number of static mapping leases configured for the interface.

DHCP Range The total size of the DHCP pool available for use on the interface.

If the Leases count approaches the Range value, then a larger pool may be required for the interface. Static mappings exist outside the range, so they do not factor into the amount of leases consumed in the pool.

Cellular

On select 3G/4G devices, the firewall is able to collect signal strength data for the Cellular graph. The signal strength is the only value plotted on the graph.

VPN Users

The VPN Users category shows the number of OpenVPN users logged in concurrently for each individual OpenVPN server.

22.2.2 Traffic Graphs

Real time traffic graphs drawn with JavaScript using NVD3 are available which update continually. These graphs can be viewed at Status > Traffic Graph, and an example of the graph can be found in Figure *Example LAN Graph*.

These traffic graphs show interface traffic as it happens, and give a clear view of what is happening "now" rather than relying on averaged data from the RRD graphs which are better for long-term views.

Only one interface is visible at a time, and this interface can be changed using the Interface drop-down list. Once an interface is chosen, the page will automatically refresh and start displaying the new graph.

Similar style traffic graphs can also be viewed on the Dashboard by adding the Traffic Graphs widget. Using the widget, multiple traffic graphs can be displayed simultaneously.

See also:

For more about the Dashboard, see *Dashboard*.

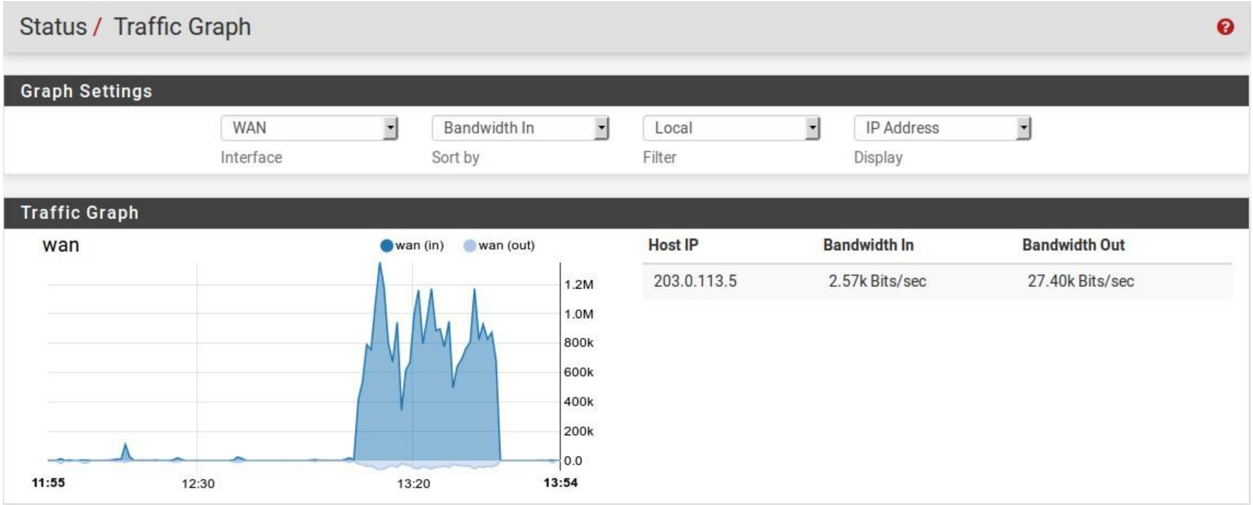

Status / Traffic Graph

Graph Settings

Interface: WAN
Sort by: Bandwidth In
Filter: Local
Display: IP Address

Traffic Graph

Host IP	Bandwidth In	Bandwidth Out
203.0.113.5	2.57k Bits/sec	27.40k Bits/sec

Fig. 16: Example LAN Graph

A table containing momentary glimpses of data being transferring from specific IP addresses is also displayed next to the traffic graph. These are limited to only displaying briefly, so ongoing transfers are more likely to show up than quick connections. Also, only connection from within that interface's primary subnet will be shown.

The display of the graph and table can be controlled using the following options:

Interface The firewall interface to use as the traffic source for the graph and the table.

Sort By Selects the sort order of the graph, either *Bandwidth In* or *Bandwidth Out*.

Filter Selects which type of hosts to display in the table

Local Shows only IP addresses within the interface network

Remote Shows only IP addresses that are not within the interface network

All Shows all IP addresses, inside and outside the interface network

Display Controls the display of the Host IP column using one of the following choices:

IP Address The IP address of the host.

Host Name The short hostname that corresponds to the IP address, as listed in DHCP static mappings, DNS Resolver host overrides, or DNS Forwarder host overrides.

Description The description that corresponds to the IP address, as listed in DHCP static mappings, DNS Resolver host overrides, or DNS Forwarder host overrides.

FQDN The fully qualified domain name that corresponds to the IP address, as listed in DHCP static mappings, DNS Resolver host overrides, or DNS Forwarder host overrides.

22.2.3 Monitoring Bandwidth Usage

With AZTCO-FW® software, there are several methods for monitoring bandwidth usage, with different levels of granularity.

pftop

If a connection is currently active, connect to the AZTCO-FW router's console (physical access or ssh) and watch the traffic flow with pftop (Option 9).

The output can be changed to show several views (press 0-8 or v to cycle) and may be sorted in various ways. Press ? for a list of available command keys while running pftop.

iftop

Install iftop from the *Package List*, then tun it from the shell (console or SSH) as follows:

```
iftop -nNpPi em0
```

Change em0 to be the interface that should be monitored.

In the above example, -nNpP tells iftop to not resolve hostnames (n) or port numbers (N), and to run in promiscuous mode (p) and also display ports in the output (P).

Press t to cycle through various views.

Darkstat

Darkstat is also available in System > Packages. Once installed, it appears under Diagnostics > darkstat. It also offers bandwidth graphs for an interface, as well as traffic to/from specific IP addresses.

NTOPNG

If even more detail is required, the ntopng package, which can also be found under System > Packages, can help. It can break down detail by IP, protocol, and so on. Once installed, it appears under Diagnostics > ntopng. It will even track where connections were made by local PCs, and how much bandwidth was used on individual connections.

The older ntop package has been replaced by ntopng.

Due to the disk resource requirements of ntop and ntopng, it is not recommended for systems that have low CPU or RAM.

Monitoring on Multiple Interfaces

Currently, darkstat and bandwidthd do not listen on multiple interfaces. ntopng will listen on multiple interfaces.

22.3 Logs

Logs in AZTCO-FW software contain recent events and messages from daemons. These messages can be stored locally on a limited basis, or forwarded to a central logging server for long-term storage, better reporting, alerting, and so on.

22.3.1 System Logs

AZTCO-FW® software logs a lot of data by default, but does so in a manner that will not overflow the storage on the firewall. The logs can be viewed in the GUI under Status > System Logs and under /var/log/ on the file system.

Some components such as DHCP and IPsec generate enough logs that they have their own logging tabs to reduce clutter in the main system log and to ease troubleshooting for these individual services. To view other logs, click the tab for the subsystem to view. Certain areas, such as System, and VPN, have sub-tabs with additional related options.

AZTCO-FW logs are contained in a binary circular log format called clog. These files are a fixed size and never grow. As a consequence of this, the log will only hold a certain amount of entries and the old entries are continually pushed out of the log as new entries are added. If log retention is an issue for an organization, the logs can be copied to another server with syslog where they may be permanently retained or rotated with less frequency. See *Remote Logging with Syslog* later in this chapter for information about syslog.

On normal installations where logs are kept on disk, they are retained across reboots. When /var is in a RAM disk, the system attempts to backup the logs at shutdown and restore them when booting. If the system does not shut down cleanly, the logs will reset.

Viewing System Logs

The system logs can be found under Status > System Logs, on the System tab. This will include log entries generated by the host itself in addition to those created by services and packages which do not have their logs redirected to other tabs/log files.

As shown by the example entries in Figure *Example System Log Entries*, there are log entries from several different areas in the main system log. Many other subsystems will log here, but most will not overload the logs at any one time. Typically if a service has many log entries it will be moved to its own tab and log file.

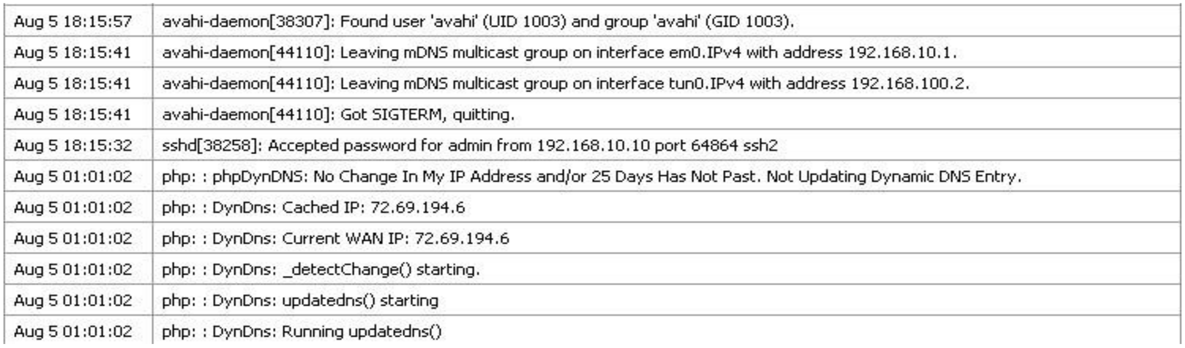

```
Aug 5 18:15:57   avahi-daemon[38307]: Found user 'avahi' (UID 1003) and group 'avahi' (GID 1003).
Aug 5 18:15:41   avahi-daemon[44110]: Leaving mDNS multicast group on interface em0.IPv4 with address 192.168.10.1.
Aug 5 18:15:41   avahi-daemon[44110]: Leaving mDNS multicast group on interface tun0.IPv4 with address 192.168.100.2.
Aug 5 18:15:41   avahi-daemon[44110]: Got SIGTERM, quitting.
Aug 5 18:15:32   sshd[38258]: Accepted password for admin from 192.168.10.10 port 64864 ssh2
Aug 5 01:01:02   php: : phpDynDNS: No Change In My IP Address and/or 25 Days Has Not Past. Not Updating Dynamic DNS Entry.
Aug 5 01:01:02   php: : DynDns: Cached IP: 72.69.194.6
Aug 5 01:01:02   php: : DynDns: Current WAN IP: 72.69.194.6
Aug 5 01:01:02   php: : DynDns: _detectChange() starting.
Aug 5 01:01:02   php: : DynDns: updatedns() starting
Aug 5 01:01:02   php: : DynDns: Running updatedns()
```

Fig. 17: Example System Log Entries

Filtering Log Entries

Every log can be searched and filtered to find entries matching a specified pattern. This is very useful for tracking down log messages from a specific service or log entries containing a specific username, IP address, and so on.

To search for log entries:

- Navigate to Status > System Logs and then the tab for the log to search
- Click ▼ in the breadcrumb bar to open the Advanced Log Filter panel
- Enter the search criteria, for example, place some text or a regular expression in the Message field
- Click ▼ Apply Filter

The filtering fields vary by log tab, but may include:

Message The body of the log message itself. A word or phrase may be entered to match exactly, or use Regular Expressions to match complex patterns.

Time The timestamp of the log message. Uses month names abbreviated to three letters.

Process The *name* of the process or daemon generating the log messages, such as sshd or check_reload_status.

PID The process ID number of a running command or daemon. In cases where there are multiple copies of a daemon running, such as openvpn, use this field to isolate messages from a single instance.

Quantity The number of matches to return in filter results. Setting this value higher than the number of log entries in the log file will have no effect, but setting it higher than the current display value will temporarily show more log messages.

The Firewall log tab has a different set of filtering fields:

Source IP Address The source IP address listed in the log entry.

Destination IP Address The destination IP address listed in the log entry.

Pass Check this option to only match log entries that passed traffic.

Block Check this option to only match log entries that blocked traffic.

Interface The friendly description name of the interface to match (e.g. WAN, LAN, OPT2, DMZ)

Source Port The source port of the log entry to match, if the protocol uses ports.

Destination Port The destination port of the log entry to match, if the protocol uses ports.

Protocol The protocol to match, such as TCP, UDP, or ICMP.

Protocol Flags For TCP, this field matches the TCP flags on the log entry, such as SA (SYN+ACK) or FA (FIN+ACK)

The filter pane is hidden by default but it can be included on the page at all times by checking Log Filter under System > General Setup.

22.3.2 Log Settings

Log settings on AZTCO-FW® software may be adjusted in two different ways:

- Globally at Status > System Logs on the Settings tab
- On each log tab where settings can override the global defaults

  To change these settings click 🔧 in the breadcrumb bar while viewing a log.

Each of these methods will be explained in detail in this section.

The global options area contains more options than the per-log settings. Only differences will be covered in detail for the per-log settings.

Global Log Settings

The global log options under Status > System Logs on the Settings tab include:

In the GUI, the Settings tab under Status > System Logs controls how the logging system behaves.

Log Message Format The format of messages logged by the system log daemon (syslogd) for local and remote logs. Both formats are handled the same way locally, but remote syslog servers may prefer one format or the other. Check the documentation of the syslog server for details.

BSD (RFC 3164, default) The default log format used by previous versions of AZTCO-FW software and natively used by FreeBSD.

syslog (RFC 5424, with RFC 3339 microsecond-precision timestamps A modern syslog message format with more precise timestamps. Also includes the hostname.

Forward/Reverse Display By default the logs are displayed in their natural order with the oldest entries at the top and the newest entries at the bottom. Some administrators prefer to see the newest entries at the top, which can be accomplished by checking this box to flip the order.

GUI Log Entries The number of log entries to display in the log tabs of the GUI by default. This does not limit the number of entries in the file, only what is shown on the page at the time. The default value is 50. The actual log files may contain much more than the number of lines to display, depending on the Log File Size.

Log File Size (Bytes) (<= 2.4.x) The size of the *circular log file*. The size of the file directly controls how many entries it can contain. The default log size is approximately 500,000 bytes (500KB).

There are roughly 20 log files, so any increase in file size will result in 20 times larger total disk utilization from logs. The current total log size and remaining disk space are displayed for reference. At the default size, the logs will hold about 2500 entries on average but it may be significantly more or less depending on the size of individual log entries.

> Warning: The new log size will not take effect until a log is cleared or reinitialized. This may be done individually from each log tab or it can be done for all logs using the 🗑 Reset Log Files button on this page. See *Adjusting the Size of Log Files* for more.

Log Packets from Default Block Rules Checked by default. When enabled, the default deny rule, which blocks traffic not matched by other rules, will log entries to the firewall log. Typically these log entries are beneficial, but in certain rare use cases they may produce undesirable log entries that are made redundant by custom block rules with logging enabled.

Log Packets from Default Pass Rules Unchecked by default. When set, logging will occur for packets matching the default pass out rules on interfaces. Setting this option will generate a large amount of log data for connections outbound from the firewall. The best practice is to only enable this for brief periods of time while performing troubleshooting or diagnostics.

Log Packets from Block Bogon Networks Rules Checked by default. When checked, if an interface has Block Bogon Networks active, packets matching that rule will be logged. Uncheck to disable the logging.

Log Packets from Block Private Networks Rules Checked by default. When checked, if an interface has Block Private Networks active, packets matching that rule will be logged. Uncheck to disable the logging.

Web Server Log When checked, log messages from the Web GUI process, nginx, will be placed in the main system log. On occasion, especially with Captive Portal active, these messages can be frequent

but irrelevant and clutter the log contents.

Raw Logs When checked, this setting disables log parsing, displaying the raw contents of the logs instead. The raw logs contain more detail, but they are much more difficult to read. For many logs it also stops the GUI from showing separate columns for the process and PID, leaving all of that information contained in the Message column.

Show Rule Descriptions Controls if, and where, the firewall log display will show descriptions for the rules that triggered entries. Displaying the rule descriptions causes extra processing overhead that can slow down the log display, especially in cases where the view is set to show a large number of entries.

> Don’t load descriptions When selected this choice will not display any rule descriptions. The description may still be viewed by clicking the action column icon in the firewall
>
> log view (e.g. ✖ or ▶).
>
> Display as column The default for new installations. Adds the rule description in a separate column. This works best if the descriptions are short, or the display is wide.
>
> Display as second row Adds a second row to each firewall log entry containing the rule description. This choice is better for long rule descriptions or narrow displays.

Tip: If the firewall logs display slowly with rule descriptions enabled, select *Don’t load descriptions* for faster performance.

Local Logging When checked, local logs are not retained. They are not written to disk nor are they kept in memory. While this saves on disk writes, it necessitates the use of remote logging so that information is not lost. This is not a best practice, as having local logs is vital for the vast majority of use cases.

Reset Log Files This button clears the data from all log files and reinitialize them as new, empty logs. This must be done after changing the log file sizes (2.4.x), and can also be used to clear out irrelevant/old information from logs if necessary.

> Warning: Resetting the log files will not save the other options on the page. If options on this page have been changed, click Save before attempting to reset the log files.

Click Save to store the new settings. The remaining options on this screen are discussed in *Remote Logging with Syslog*.

Log Rotation Settings

Starting with AZTCO-FW software version 2.5.0, the system logs are kept in a plain text format and periodically rotated. The options in this section control how the system handles log rotation.

> Note: The options in this section of the page are global only, and cannot be changed for individual logs.

Log Rotation Size (Bytes) This field controls the size at which the system will rotate logs. The default size is 500 KiB per log file. There are nearly 20 log files, so plan space accordingly.

This does also does not account for space used by rotated log files.

> Note: Increasing this value allows every log file to grow to the specified size, so disk usage may increase significantly. Log file sizes are checked once per minute to determine if rotation is necessary, so a very rapidly growing log file may exceed this value.

Log Compression The type of compression to use when the system rotates log files. Compressing rotated log files saves disk space, and the compressed logs remain available for display and searching in the GUI. Though processing large compressed files can be time consuming, most use cases will not notice significant slowness.

The types of compression available are bzip2 (default), gzip, xz, zstd, and none (disables compression). All of the options which use compression are reasonably fast and offer good compression rates. Some may compress better than others, others are slightly faster, but ultimately the decision is up to the environment and the administrators.

> Warning: The type of compression used by all log files must be identical. When changing this value, the system must remove all previously rotated compressed log files.

Log Retention Count The number of rotated log files to keep before the oldest copy is removed. Keeping more log files will consume more disk space, but compressed logs files do not consume nearly as much space as decompressed logs.

Per-Log Settings

To change per-log settings, visit the log tab to change and then click 🔧 in the breadcrumb bar to expand the settings panel.

On this panel, several options are displayed. Most of the options will show the global default value or have a General Logging Options Settings choice which will use the global value and not the per-log value.

The per-log settings panel for each tab only displays options relevant to that log. For example, the options to log default block or pass rules are displayed only when viewing the Firewall log tab.

Each per-log settings panel has at least the following options: Forward/Reverse Display, GUI Log Entries, Log File Size (Bytes) (2.4.x), and Formatted/Raw Display. For each of these, a value which will only apply to this log may be set. For more information on how these options work, see *Global Log Settings* above.

Click Save to store the new log settings.

Note: If the log file size was changed, after saving, open the settings panel again and click the 🗑 Clear Log button to reset the log using the new size.

See also:

- *Remote Logging with Syslog*
- *Accessing Firewall Services over IPsec VPNs*
- *Working with Log Files*
- *Adjusting the Size of Log Files*

22.3.3 Remote Logging with Syslog

The Remote Logging options under Status > System Logs on the Settings tab allow syslog to copy log entries to a remote server.

The logs kept by AZTCO-FW® software on the firewall itself are of a finite size. Copying these entries to a syslog server can aid troubleshooting and allow for long-term monitoring. Having a remote copy can also help diagnose events that occur before a firewall restarts or after they would have otherwise been lost due to clearing of the logs or when older entries are cycled out of the log, and in cases when local storage has failed but the network remains active.

Warning: Corporate or local legislative policies may dictate the length of time logs must be retained from firewalls and similar devices. If an organization requires long-term log retention for their own or government purposes, a remote syslog server is required to receive and retain these logs.

Warning: Logs sent using this method are delivered in the clear (not encrypted) unless the logs are sent through a VPN or using a mechanism such as *Stunnel package*. As an alternative, consider using the syslog-ng package which supports encrypted syslog.

The following options are available for remote logging:

Source Address Controls where the syslog daemon binds for sending out messages. In most cases, the default (*Any*) is the best option, so the firewall will use the address nearest the target. If the destination server is across a tunnel mode IPsec VPN, however, choosing an interface or Virtual IP address inside the local Phase 2 network will allow the log messages to flow properly over a tunnel.

IP Protocol When choosing an interface for the Source Address, this option gives the syslog daemon a preference for either using IPv4 or IPv6, depending on which is available. If there is no matching address for the selected type, the other type is used instead.

Remote Log Servers Enter up to three remote servers using the boxes contained in this section. Each remote server can use either an IP address or hostname, and an optional UDP port number. If the port is not specified, the default *syslogd* port, 514, is assumed.

A syslog server is typically a server that is directly reachable from the firewall on a local interface. Logging can also be sent to a server across a VPN.

Warning: Do not send log data directly across any WAN connection or unencrypted site-to-site link, as it is plain text and could contain sensitive information.

Note: The syslog daemon only supports sending messages over UDP. To send syslog messages over TCP, consider using the syslog-ng package.

Remote Syslog Contents The options in this section control which log messages will be sent to the remote log server.

Everything When set, all log messages from all areas are sent to the server.

System Events Main system log messages that do not fall into other categories.

Firewall Events Firewall log messages in raw format. The format of the raw log is covered in *Raw Filter Log Format*.

DNS Events Messages from the DNS Resolver (unbound), DNS Forwarder (dnsmasq), and from the filterdns daemon which periodically resolves hostnames in aliases.

DHCP Events Messages from the IPv4 and IPv6 DHCP daemons, relay agents, and clients.

PPP Events Messages from PPP WAN clients (PPPoE, L2TP, PPTP)

General Authentication Events Log messages about authentication events, such as for the GUI or certain types of VPNs.

Captive Portal Events Messages from the Captive Portal system, typically authentication messages and errors.

VPN Events Messages from VPN daemons such as IPsec and OpenVPN, as well as the L2TP server and PPPoE server.

Gateway Monitor Events Messages from the gateway monitoring daemon, dpinger

Routing Daemon Events Routing-related messages such as UPnP/NAT-PMP, IPv6 routing advertisements, and routing daemons from packages like OSPF, BGP, and RIP.

Network Time Protocol Events Messages from the NTP daemon and client.

Wireless Events Messages from the Wireless AP daemon, hostapd.

To start logging remotely:

- Navigate to Status > System Logs on the Settings tab
- Check Send log messages to remote syslog server
- Configure the options as described above
- Click Save to store the changes.

If a syslog server is not already available, it is fairly easy to set one up. Almost any UNIX or UNIX-like system can be used as a syslog server. FreeBSD is described in the following section, but others may be similar.

Setup Syslog on the Logging Host

FreeBSD

First, configure the syslog server to accept remote connections which means running it with the -a <subnet> or similar flag.

On FreeBSD, edit /etc/rc.conf and add this line:

```
syslogd_flags=" -a 192.168.1.1 "
```

Where 192.168.1.1 is the IP address of the AZTCO-FW firewall.

More complex allow rules for syslog are also possible, like so:

```
syslogd_flags=" -a 10.0.10.0/24:*"
```

Using that parameter, syslog will accept from any IP address in the 10.0.10.0 subnet (mask 255.255.255.0) and the messages may come from any UDP port.

Now, edit /etc/syslog.conf and add a block at the bottom:

```
!*
+*

+AZTCO-FW
*.*                          /var/log/AZTCO-FW.log
```

Where AZTCO-FW is the hostname of the AZTCO-FW firewall. An entry may also need to be added in /etc/hosts for that system, depending on the DNS setup. Logs may be split separate files. Use the /etc/syslog.conf file on the AZTCO-FW firewall for more details on which logging facilities are used for specific items.

```
192.168.1.1                  AZTCO-FW      AZTCO-FW.example.com
```

The log file may also need to be created manually with proper permissions:

```
touch /var/log/AZTCO-FW.log chmod 640
/var/log/AZTCO-FW.log
```

Now restart syslog:

```
/etc/rc.d/syslogd restart
```

Windows

Setting this up on Windows entirely depends on which syslog server is being used. Consult the documentation for more information on configuration.

There is a free multi-purpose utility that can act as a syslog server, which can be found here: http://tftpd32.jounin.net/

Kiwi Syslog Server is free for up to 5 devices. http://www.kiwisyslog.com/downloads.aspx Linux

Configuration of the system logger on Linux depends on the distribution. Consult the distribution's documentation on how to change the behavior of syslogd. It should be similar in many cases to the alterations in the FreeBSD section.

OpenBSD

The configuration for OpenBSD is similar to FreeBSD, with the following notes:

1. The option to accept remote syslog events is -u.
2. This option may be enabled using *rcctl(8)*:

```
rcctl set syslogd flags -u
```

1. To restart the syslogd service:

```
rcctl restart syslogd
```

Other Logging Servers

Other log systems such as Splunk, ELSA, or ELK may also be used but the methods for implementing them are beyond the scope of this document. If such a system is syslog-compatible, then the AZTCO-FW software side should be fairly simple to setup as it would be for any other syslog system.

22.3.4 Adjusting the Size of Log Files

AZTCO-FW® software stores its logs in *binary circular log files* that never grow in size. These are known as "clog" files. The fixed size prevents the logs from filling up available storage space and eliminates the need for rotation, but the down side is that the logs wrap around once full, losing older messages in the process. These log files are held in /var/log which may optionally be a RAM disk.

The Log File Size (Bytes) field in the *Log Settings* at Status > System Logs, Settings tab, allows the user to specify the amount of storage to allocate per log file.

Short Version

To change the log file sizes:

- Navigate to Status > System Logs, Settings tab
- Enter a new value in Log File Size (Bytes), being careful not to overfill the disk containing the logs.
- Click Save
- Click Reset Log Files

Details

The default size for a log file is *511488* (~500KB) which can generally hold between 2000-3000 log entries but varies by entry size. The time span covered by logs depends entirely on how much data is logged. A quiet log file could contain months or even years of information, a busy log file may only contain minutes.

Underneath the text for Log File Size (Bytes) the current and available disk space is displayed based on the current log file sizes and their location. For example:

```
Disk space currently used by log files: 9.8M. Remaining disk space for log files: 11G
```

There are approximately 20 log files affected by the size control. The value entered in Log File Size (Bytes) is for a single log file, so the actual usage will be approximately 20 times that value. As shown above, with the default value of 511488 (~500KB), the firewall uses nearly 10MB of total log space. If the size of the logs is increased to *1024000* (1MB), then nearly 20MB would be used for logs. Be certain before changing the Log File Size (Bytes) value that the disk has enough space to hold all of the log files.

The change to the log size only takes effect when log files are reset. Which only happens when the logs are manually reset. Log files may be reset individually using the Clear button on the various log tabs, or by using the master Reset Log Files button on the *Log Settings* page.

22.3.5 Viewing the Firewall Log

The firewall creates log entries for each rule configured to log and for the default deny rule. There are several ways to view these log entries, each with varying levels of detail. There is no clear “best” method since it depends on the preferences and skill level of the firewall administrators, though using the GUI is the easiest method.

Tip: The logging behavior of the default deny rules and other internal rules can be controlled using the Settings tab under Status > System Logs. See *Log Settings* for details.

Like other logs, the firewall log only retains a certain number of entries. If the needs of an organization require a permanent record of firewall logs for a longer period of time, see *Remote Logging with Syslog* for information on copying these log entries to a syslog server as they happen.

Viewing in the WebGUI

The firewall logs are visible in the WebGUI at Status > System Logs, on the Firewall tab. From there, the logs can be viewed as a parsed log, which is easier to read, or as a raw log, which contains more detail. There is also a setting to show these entries in forward or reverse order. If the order the log entries being displayed is unknown, check the timestamp of the first and last lines, or check *Log Settings* for information on how to view and change these settings.

The parsed WebGUI logs, seen in Figure *Example Log Entries Viewed From The WebGUI*, are in 6 columns: Action, Time, Interface, Source, Destination, and Protocol.

Action Shows what happened to the packet which generated the log entry (e.g. pass or block)

> The Action icon is a link which, when clicked, looks up and displays the rule which caused the log entry. More often than not, this says “Default Deny Rule”, but when troubleshooting rule issues it can help narrow down suspects.

Time The time that the packet arrived.

Interface Where the packet entered the firewall.

The GUI prints a I character next to the interface if a rule matched a packet in the *outbound* direction. The vast majority of rules match in the inbound direction, so the direction is omitted in that case.

Rule The firewall rule description and ID number which generated the log entry, if available. This column only appears when rule descriptions are set to appear in a separate column. They may also be shown in a separate row, or disabled entirely. See *Log Settings* for details.

Source The source IP address and port.

The i icon next to the source and destination IP addresses, when clicked, makes the firewall perform a DNS lookup on the IP address. If the address has a valid hostname it will be displayed underneath the IP address in all instances of that address on the page.

The ⊟ icon next to the source IP address and the ⊞ icon next to the destination IP address are for adding firewall rules with EasyRule. See *Using EasyRule to Add Firewall Rules* for details.

Destination The destination IP address and port.

Protocol The protocol of the packet, e.g. ICMP, TCP, UDP, etc.

Log entries for TCP packets have extra information appended to the protocol field displaying TCP flags present in the packet. These flags indicate various connection states or packet attributes. The meanings for each flag are outlined in *TCP Flags*.

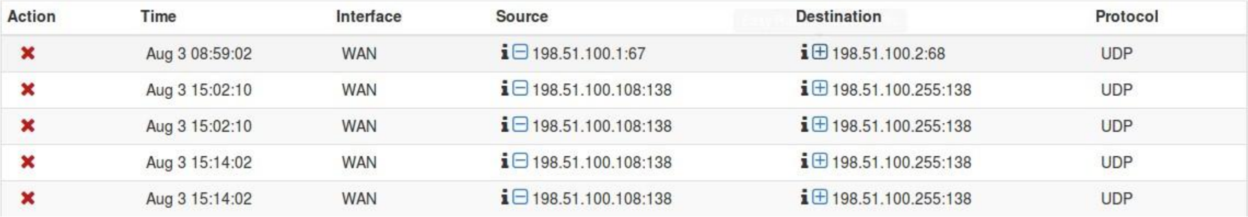

Action	Time	Interface	Source	Destination	Protocol
✖	Aug 3 08:59:02	WAN	i ⊟ 198.51.100.1:67	i ⊞ 198.51.100.2:68	UDP
✖	Aug 3 15:02:10	WAN	i ⊟ 198.51.100.108:138	i ⊞ 198.51.100.255:138	UDP
✖	Aug 3 15:02:10	WAN	i ⊟ 198.51.100.108:138	i ⊞ 198.51.100.255:138	UDP
✖	Aug 3 15:14:02	WAN	i ⊟ 198.51.100.108:138	i ⊞ 198.51.100.255:138	UDP
✖	Aug 3 15:14:02	WAN	i ⊟ 198.51.100.108:138	i ⊞ 198.51.100.255:138	UDP

Fig. 18: Example Log Entries Viewed From The WebGUI

The GUI can also filter log output to find specific entries, so long as they exist in the current log. Click ▼ to display the filtering options. See *Filtering Log Entries* for more information.

Viewing from the Console Menu

Option 10 from the console menu views and follows the filter.log in real time. An easy example is a log entry like that seen above in Figure *Example Log Entries Viewed From The WebGUI*:

```
Aug 3 08:59:02 master filterlog: 5,16777216,,1000000103,igb1,match,block,in,4,0x10,,
↪128,0,0, none,17,udp,328,198.51.100.1,198.51.100.2,67,68,308
```

This single line shows that the log entry was triggered by rule id 1000000103, which resulted in a block action on the igb1 interface. The source and destination IP addresses are shown near the end of the log entry, followed by the source and destination port. Packets from other protocols may show significantly more data.

Viewing from the Shell

When using the shell, either from SSH or from the console, there are numerous options available to view the filter logs.

When directly viewing the contents of the log file, the log entries can be quite complex and verbose.

For information on viewing logs from the shell, see *Working with Log Files*.

Viewing parsed log output in the shell

There is a simple log parser written in PHP which can be used from the shell to produce reduced output instead of the full raw log. To view the parsed contents of the current log, run:

2.5.0 and later:

```
# cat /var/log/filter.log | filterparser.php
```

Older versions:

```
# clog /var/log/filter.log | filterparser.php
```

The script prints the log entries one per line, with simplified output:

```
Aug 3 08:59:02 block igb1 UDP 198.51.100.1:67 198.51.100.2:68
```

Finding the rule which caused a log entry

When viewing one of the raw log formats, the log includes the rule ID number for an entry. This rule number can be used to find the rule which caused the match. The following example locates the rule with id 1000000103:

```
# pfctl -vvsr | grep 1000000103 @5(1000000103) block drop in log inet all label "Default deny rule IPv4"
```

As shown in the above output, this was the default deny rule for IPv4.

22.3.6 Gateway Logs

The gateway logs can be found through the AZTCO-FW® webGUI under Status > System Logs on the System/Gateways sub-tab.

This log contains entries from the gateway monitoring daemon, *dpinger*, which can generate a significant amount of logging with many gateways to monitor.

The entries found here will record events such as when a gateway is down, or in an alarm state, or has returned to an online state.

22.3.7 NTP Logs

The NTP daemon Log view will show any logs generated by the *Network Time Protocol daemon* and logs from the NTP client *ntpdate* that performs large time skew updates.

22.3.8 Package Logs

In AZTCO-FW® software, packages which support logging to this central location will have their logs displayed here. If a package's logs do not show up, contact the package maintainer to see if the package can be updated to support this mechanism.

Some packages will log to the main system log or a related tab inside the system logs (Status > System Logs). Others may keep their own logs in a separate location. Some packages, such as Squid and Snort, offer configuration options to control where and how logs are made. Some logs may need to be viewed outside the webGUI or via Diagnostics > Command.

22.3.9 PPP Logs

The PPP logs tab displays any events from the PPP system for WAN type connections, not locally-hosted servers. This would be for WANs that connect using PPPoE, 3G networks or in some rare cases, dialup, etc.

22.3.10 Routing Logs

The Routing logs are located at Status > System Logs on the System/Routing tab.

This log contains entries from routing-related processes for both IPv4 and IPv6, including:

- radvd (IPv6 Router Advertisements)
- routed (RIP)
- olsrd
- zebra (Quagga)
- ospfd (Quagga or OpenOSPFD)
- bgpd (OpenBGPD)
- miniupnpd (UPnP/NAT-PMP)

22.3.11 IPsec Logs

The IPsec logs show output from the IPsec daemon, handled by strongswan. Normal output, successful connections, as well as errors are all displayed here.

Where possible, if a log message contains an IP address of a configured IPsec tunnel, that tunnel's description is prepended to the log entry.

Entries found in these logs are covered in depth in *IPsec Logs*, and some errors are covered in *Troubleshooting IPsec VPNs*.

22.3.12 OpenVPN Logs

The OpenVPN logs found through the AZTCO-FW® webGUI at Status > System Logs and the OpenVPN tab show output from the OpenVPN daemon(s) in use, both clients and servers. Messages are shown in the logs for successful connections as well as failures and errors.

If there are no log entries for a server after the process starts, traffic likely is not reaching the OpenVPN daemon. Check the WAN-side firewall rules and the address/port used by the client.

There are several OpenVPN troubleshooting articles found in *Troubleshooting*.

22.3.13 Captive Portal Authentication Logs

The Captive Portal Authentication Logs are available through the AZTCO-FW® webGUI at Status > System Logs, on the Portal Auth tab. The logs list login information from the Captive Portal system. Squid authentication may also be displayed in this tab.

22.3.14 Wireless Logs

The Wireless logs can be found in the AZTCO-FW® webGUI under Status > System Logs on the System/Wireless tab.

These logs contain entries from the *hostapd* daemon which handles wireless access point connections. This process can be overly verbose when handling client traffic, logging rekeys and other information that can otherwise clutter the main system log.

22.3.15 L2TP Logs

A record of login and logout events is kept on Status > System Logs, on the VPN tab, under L2TP Logins.

Each login and logout is recorded with a timestamp and username, and each login will also show the IP address assigned to the L2TP client. The full log can be found on the L2TP Raw tab.

22.3.16 DHCP Logs

The DHCP log view at Status > System Logs on the DHCP Tab, displays messages and events from the DHCP Daemon and the DHCP client for WANs.

Each DHCP request and reply from DHCP clients is shown here, along with events and errors. IP addresses, MAC addresses, and client-supplied hostnames are all visible in the logs.

See also:

- *Troubleshooting “login on console as root” Log Messages*
- *Troubleshooting “promiscuous mode enabled” Log Messages*
- *Troubleshooting ARP Move Log Messages*

CHAPTER

TWENTYTHREE

DIAGNOSTICS

These documents cover functions found under the Diagnostics menu in AZTCO-FW® software.

23.1 DNS Lookup

Diagnostics > DNS Lookup performs simple forward and reverse DNS queries to obtain information about an IP address or hostname, and also to test the DNS servers used by the firewall.

To perform a DNS Lookup:

- Navigate to Diagnostics > DNS Lookup
- Enter a Hostname or IP address to query
- Click 🔍 Lookup

The page displays the results of the DNS query along with supporting information and options.

23.1.1 Results

The Results panel contains addresses returned by the DNS query along with the record type.

Underneath the results is a table containing the resolution Timings per server. This shows how fast each of the configured DNS servers responded to the specified query, or if they never responded.

The More Information panel contains links to ping and traceroute functions on the firewall for this host.

23.1.2 Aliases

When performing a DNS lookup, the GUI can also create a *firewall alias* from the results of the query. The name of the alias is the text entered for the DNS query but with . characters replaced by _. For example, a DNS lookup for example.com results in an alias named example_com.

Click ✚ Add alias to create an alias containing the results of the query.

If an alias already exists with that name, the button is labeled Update Alias instead. That version of the button will replace the contents of the existing alias with the current results of the DNS lookup.

23.2 Editing Files on the Firewall

Diagnostics > Edit File contains a file editor that allows editing and creating files on the AZTCO-FW® filesystem.

Enter the filename to edit or create in Save / Load from path, or click Browse and locate the file. Once the path is filled in, press Load. If a new file is being created, type the path and filename to which it will be saved.

Edit or create the text, then press Save when finished. For existing files, the contents will be saved. For new files, the file will be created.

> Warning: Be careful when choosing a file to edit! It is very easy to edit the wrong file, or break a piece of code, and render the system unusable. Use of this tool is not recommended except under guidance of support or when there is sufficient knowledge to use it without causing unintended side effects.

23.3 Command Prompt

The command prompt, available at Diagnostics > Command Prompt, executes shell commands, PHP code, and can download or upload whole files.

> Warning: Exercise caution using any of these utilities. Executing commands and PHP code improperly can render the firewall unusable. Use of this tool is not recommended except under the guidance of a support representative or if there is sufficient knowledge on the part of the user.

23.3.1 Execute Shell Commands

To execute a shell command:

- Navigate to Diagnostics > Command Prompt
- Enter the command into the Command box under Execute Shell command
- Click Execute

Commands are executed as if they were run from a console command line, and the page prints the results when the command terminates.

> Warning: Commands must run and then stop or return.
>
> Commands that run indefinitely, such as ping without a count or tcpdump without a limit set will never stop or return output, and will be left running indefinitely in the background until they are manually killed.
>
> Interactive commands, such as vi will fail similarly, or may exit due to other issues with the terminal being non-interactive.

Previously used commands from this session can be recalled with the 🔍 and » buttons. The browser will forget the previous command list once it leaves the page.

23.3.2 Download

To download a file from the firewall filesystem:

- Navigate to Diagnostics > Command Prompt
- Enter the full path name in File to download
- Click Download

23.3.3 Upload

To upload a file:

- Navigate to Diagnostics > Command Prompt
- Click Browse
- Locate and select the file on the local client computer
- Click Upload

Note: Uploaded files are placed in /tmp/ and can then be moved to alternate locations by other functions (such as the Execute Shell Command feature).

23.3.4 PHP Execute

This page can also execute PHP code.

- Navigate to Diagnostics > Command Prompt
- Type or paste PHP code into the Execute PHP Commands text area
- Click Execute

The GUI displays the output from the PHP code above the text area, or an error if the it could not run the code.

23.4 Ping Host

The firewall can send ICMP echo reqests, also known as “pings”, to hosts over the network to test if they respond, and to measure latency between the firewall and target hosts. A basic ping test can be performed at the console, and a more detailed test is available in the GUI at Diagnostics > Ping.

23.4.1 Ping Options

When performing a ping test from the GUI, the following options are available:

Hostname A hostname or IP address to which the firewall will send pings.

IP Protocol The address type to ping when a hostname is entered that has both A (*IPv4*) and AAAA (*IPv6*) records.

Source Address The IP address from which the ping will be sent. This is especially important when testing LAN-to-LAN VPN connectivity. The default choice allows the operating system to automatically select the closest address to the target, based on the routing table.

Maximum Number of Pings The number of ping requests the firewall will send during this test. A higher count will take longer to complete and display results, especially if the target is down. Default is 3.

Seconds Between Pings The number of seconds to wait between sending ping requests. Default is 1 second.

23.4.2 Ping from the GUI

To perform a ping test from the GUI:

- Navigate to Diagnostics > Ping
- Fill in the *Ping Options*

Note: At a minimum the Hostname is required.

- Click Ping to start the test
- Wait for the GUI to display the test results

The GUI will display the results of the test automatically once complete. Do not navigate away from the page while the test is running.

23.4.3 Ping from the Console

A simple ping test may also be performed at the console menu, but without the additional options mentioned earlier. See *Ping host* for more information.

- Access the console menu locally or via SSH with an admin-level account (admin, root, or another privileged account using sudo).
- Enter the menu option which corresponds with Ping Host (e.g. 7)
- Press Enter
- Enter the IP address or hostname to ping

- Press Enter to start the test
- Wait for the test to complete.

  The console outputs the test results in real time, and pauses afterward.
- Press Enter to return to the menu

23.5 Halting and Powering Off the Firewall

The firewall may be shut down safely by the Halt function available at Diagnostics > Halt System or from the console menu.

Warning: The best practice is to never cut power from a running system. Halting before removing power is always the safest choice.

Aafter the operating system halts, the device power will also be turned off if that feature is supported by the hardware.

23.5.1 Halt from the GUI

To halt the operating system from the GUI:

- Navigate to Diagnostics > Halt System
- Click ⏹ Halt
- Click OK to confirm the action and start the halt process

23.5.2 Halt from the Console

- Access the console menu locally or via SSH with an admin-level account (admin, root, or another privileged account using sudo).
- Enter the menu option which corresponds with Halt system (e.g. 6)
- Press Enter
- Enter the y to confirm the action
- Press Enter to start the halt process

23.6 Rebooting the Firewall

AZTCO-FW® software can be rebooted safely and returned to an operational state using the page at Diagnostics > Reboot System or the console.

23.6.1 Reboot Methods

The following reboot methods are possible, but available options may be limited depending on the platform and installation options.

Reboot normally Performs a normal reboot in the traditional way. This method is always available.

Reroot Performs a “reroot” style reboot, which is faster than a traditional reboot but does not restart the entire operating system. All running processes are killed, all filesystems are remounted, and then the system startup sequence is run again. This type of restart is much faster as it does not reset the hardware, reload the kernel, or need to go through the hardware detection process.

This option is not compatible with ZFS.

28.5. Halting and Powering Off the Firewall

Reboot into Single User Mode Restarts the firewall into single user mode for diagnostic purposes. The firewall cannot automatically recover from this state, console access is required to use single user mode and reboot the firewall.

This option is not compatible with ARM-based systems.

> Warning: In single user mode, the root filesystem defaults to read-only and other filesystems are not mounted. The firewall also does not have an active network connection. This option must only be used under the guidance of a support representative or a FreeBSD user with advanced knowledge.

Reboot and run a filesystem check Reboots the firewall and forces a filesystem check using fsck, run five times. This operation can often correct issues with the filesystem on the firewall.

This option is not compatible with ARM-based systems.

23.6.2 Reboot from the GUI

To reboot from the GUI:

- Navigate to Diagnostics > Reboot System
- Select the *Reboot Method*
- Click 🔧 Submit to reboot the system immediately

23.6.3 Reboot from the Console

To reboot from the console:

- Access the console menu locally or via SSH with an admin-level account (admin, root, or another privileged account using sudo).
- Enter the menu option which corresponds with Reboot system (e.g. 5)
- Press Enter

- Enter the letter which corresponds with the desired *Reboot Method*
- Press Enter

Note: The single user mode and filesystem check options require an uppercase letter to be entered to confirm the action. This is necessary to avoid activating the options accidentally. The reboot and reroot options may be entered in upper or lower case.

23.7 Testing a TCP Port

The Diagnostics > Test Port page performs a simple TCP port connection test to check if the firewall can communicate with another host. This tests if a host is up and accepting connections on a given port, at least from the perspective of the firewall.

No data is transmitted to the remote host by this test, it only attempts to open a connection and optionally displays the data sent back from the server.

In the default mode the test attempts a simple TCP handshake (SYN, SYN+ACK, ACK), and if the attempt succeeds, it reports the result.

Note: This test does not function for UDP since there is no way to reliably determine if a UDP port accepts connections in this manner.

To perform a test:

- Navigate to Diagnostics > Test Port
- Fill in the fields on the page. The Hostname and Port fields are required, the rest are optional.
- Click 🔧 Test.

The following options are available on this page:

Hostname This is the IP address or hostname of the target system. This is a required field.

Port This is the TCP port on the target used by the test. This is a required field and must be a valid port number, meaning an integer between 1 and 65535.

Source Port An optional specific source port for the query. This is unnecessary in most cases.

Remote Text If checked, this option shows the text given by the server when connecting to the port. The server is given 10 seconds to respond, and this page will display all of the text sent back by the server in those 10 seconds. As such, the test will run for a minimum of 10 seconds when performing this check.

> Note: Not all daemons will output text to the user on connect, so this may be blank even if the service is working properly. For example, an SMTP server will respond with a welcome message, as will FTP, but an HTTP daemon will not send any text.

Source Address A specific source IP address or IP Alias/CARP Virtual IP from which the query will be sent. The service being tested may require a specific source IP address, network, etc, in order to make a connection.

IP Protocol This option selects either *IPv4* or *IPv6* to control which type of IP address is used when testing a hostname. If the connection is forced to IPv4 or IPv6 and the hostname does not contain a result using that protocol, the test will produce an error. For example if forced to IPv4 and given a hostname that only returns an IPv6 IP address (AAAA record), the test will fail.

23.9 Traceroute

The traceroute page, located at Diagnostics > Traceroute, works like the *traceroute* command found on many platforms. It sends special packets which, as the name implies, trace a route across the network from the AZTCO-FW® host to a remote host. A list of hops between hosts will be displayed, along with response times, as long as the intervening hosts support (or don't filter) traffic required for traceroute to work.

- Host: A hostname or IP address to which the route will be traced.
- IP Protocol: The address type to use when a hostname is entered that has both A (*IPv4*) and AAAA (*IPv6*) records.
- Source Address: The IP address from which the trace will be sent. This is especially important when testing LAN-to-LAN VPN connectivity.
- Maximum number of hops: The maximum length of the path to trace. The trace will stop if the path cannot be traced completely after this number of hops.
- Reverse Address Lookup: When checked, traceroute will attempt to perform a PTR lookup to locate hostnames for hops along the path. Will slow down the process as it has to wait for DNS replies.
- Use ICMP: By default, traceroute uses UDP but that may be blocked by some routers. Check this box to use ICMP instead, which may succeed.

The output will be displayed once the trace is complete. The Stop button may be pressed at any time to see the current output of the trace if it is still running or stalled.

CHAPTER

TWENTYFOUR

PACKAGES

24.1ACME package

Let’s Encrypt is an open, free, and completely automated Certificate Authority from the non-profit Internet Security Research Group (ISRG). The goal of Let’s Encrypt is to encrypt the web by removing the cost barrier and some of the technical barriers that discourage server administrators and organizations from obtaining certificates for use on Internet servers, primarily web servers. Most browsers trust certificates from Let’s Encrypt. These certificates can be used for web servers (HTTPS), SMTP servers, IMAP/POP3 servers, and other similar roles which utilize the same type of certificates.

The ACME Package for AZTCO-FW® software interfaces with Let’s Encrypt to handle the certificate generation, valida- tion, and renewal processes.

Certificates from Let’s Encrypt are domain validated, and this validation ensures that the system requesting the cer- tificate has authority over the domain in question. This validation can be performed in a number of ways, such as by proving ownership of the domain’s DNS records or hosting a file on a web server for the domain.

By using a certificate from Let’s Encrypt for a web server, including a firewall running AZTCO-FW software, the browser will trust the certificate and show a green check mark, padlock, or similar indication. The connection will be encrypted without the need for a client to manually trust an invalid or self- signed certificate.

Let’s Encrypt certificates are valid for a period of 90 days, so they must be renewed periodically. The ACME package automates this renewal by using a cron job to check once per day to see if a certificate needs to be renewed.

24.1.1ACME Overview

Rate Limits

Let's Encrypt enforces rate limitations when using the production validation system, such as:

- Five validation failures per account, per hostname, per hour
- Each certificate may have at most 100 SAN entries
- Only 50 certificates may be created per domain per week

A testing validation system exists for developers who are programming clients or administrators testing their settings. The test system has higher limits, which aid testing and development, but the test system does not produce certificates which are trusted publicly.

Security Limitations

When validating using a method such as webroot or standalone the service must be available to the Internet on its standard port: 80 for HTTP or 443 for TLS-ALPN. This is a security limitation to prevent a user from running an alternate web server on a high- numbered port and obtaining a certificate for a server they do not normally control.

Validation Process

When creating a certificate, one or more fully qualified domain names (FQDNs) are listed on the certificate in the SAN list. Let's Encrypt will query each of these domain names in DNS in different ways depending on the validation method.

When a validation method starts, the client obtains an authorization value from the server (authz).

For DNS-based methods, Let's Encrypt checks for a TXT record in the form of `_acme-challenge.<domain name>` which must contain the authorization value. This proves that the person or system requesting the certificate controls DNS records for the domain.

For File-based methods such as webroot or standalone, Let's Encrypt connects to an IP address obtained by resolv- ing the A record for the FQDN and requests a file from the web server at `.well-known/acme-challenge/` underneath the webroot directory. This file contains the authorization value. This proves that the person or system requesting the certificate controls web server for the domain name.

24.1.2Obtaining a Certificate

These instructions cover the general process of obtaining a certificate. Specific settings will vary by deployment, and each section below links to the settings for each area.

Generate an Account Key

Before a certificate can be created by the firewall, the firewall must first obtain an account key. This key is typically unique for each server, but can be shared.

For users unfamiliar with Let’s Encrypt, the first key should be for the staging system which has no rate limits but is not valid for public use. Once a certificate is successfully issued by the staging system, create an account key for the production system and then issue the certificate again using that key.

To create and register an account key:

- Navigate to **Services > ACME Certificates**, **Account Keys** tab
- Click + **Add**
- Fill in the info as described in *Account Key Settings*
- Click + **Create new account key**
- Click 🔑 **Register ACME account key**
- Click **Save**

Create a certificate

The next step is to create a certificate entry.

- Navigate to **Services > ACME Certificates**, **Certificates** tab
- Click + **Add**
- Fill in the info as described in *Certificate Settings*
- Add one or more **Domain SAN List entries** (*Certificate Settings*) with appropriate validation settings (*Validation Methods*)
- Add one or more **Actions list** entries (*Certificate Settings*)
- Click **Save**

Configure General Settings

The last step is to enable at least the **Cron Entry** to ensure that the ACME package will automatically renew certificates before they expire. See *General Settings* for detailed descriptions of the options.

- Navigate to **Services > ACME Certificates**, **General Settings** tab
- Check **Cron Entry**
- Check **Write Certificates** (optional)
- Click **Save**

24.1.3ACME Package Settings

These sections describe the settings for each tab in the ACME package.

Account Key Settings

An ACME account key has the following settings:

Name A short name for the key

Description A longer string describing the key

ACME Server The ACME server to which this key will be registered by the package.

Currently supported options are:

Let’s Encrypt Staging ACMEv2 Use this server when testing the certificate validation process. Does not produce publicly trusted certificates.

Let’s Encrypt Production ACMEv2 Use this server for trusted production certificates.

BuyPass Production ACMEv2 An alternative service for ACME certificates

E-Mail Address An e-mail address which Let’s Encrypt will use to send certificate expiration notices if certificates are not renewed in a timely manner.

Account Key The RSA private key for this entry. To create a new key, click + Create new account key.

Certificate Settings

Certificate entries have the following settings:

Name A short name for the certificate

Description A longer string describing the certificate

Status Whether or not this entry is active

- **Active** This entry will be processed manually and by the Cron job (*General Settings*)
- **Disabled** This entry will be ignored

Acme Account The account key ACME will use when requesting the certificate (see *Generate an Account Key*)

Private Key The key length of the private key for this certificate. May be either RSA or ECDSA in several pre-defined sizes. Select *Custom* to manually enter a private key generated elsewhere

2048-bit RSA is an acceptable default choice, but larger keys are more secure

OCSP Must Staple When set, ACME will configure the certificate request for OCSP Stapling

> **Warning:** Do not enable this option unless all consumers of the certificate support OCSP Stapling.

Domain SAN List A list of all domain names which will be included in this certificate as Subject Alter- native Name (SAN) entries.

> **Note:** A certificate can contain up to 100 SAN entries, and they can use the same or different update methods. Each SAN must be individually validated by Let's Encrypt before a certificate will be issued.

Mode Whether or not this SAN is active in the certificate

Domain Name The domain name for a SAN entry in this certificate (e.g. `www. example.com`)

Method The method used by ACME to validate ownership of this domain. Method settings are described in (*Validation Methods*)

Click ➕ **Add** for additional SAN entries

DNS Providers also have some common settings which appear for all types:

DNS Alias An alternative domain name used by the validation process. Instead of updat- ing the DNS record for **Domain Name** directly, the package uses this domain name is used instead. See *DNS Alias Mode* for details.

DNS Alias Mode When set, controls whether or not the DNS alias mode used is Chal- lenge Alias (Unchecked, Default) or Domain Alias (Checked). See *DNS Alias Mode* for details.

DNS-Sleep The amount of time the ACME validation process will wait after making DNS changes before attempting to validate. Some DNS services take a few minutes to propagate entries after making backend changes.

The default settings are typically sufficient, but slower providers may require a longer sleep time.

Actions List Commands to run after the package renews a certificate.

Mode Whether or not this action is active

Command Full path to command and arguments, service name, or name of script

Method Defines how the **Command** is executed by the package

Shell Command The **Command** is a full path to a shell command and its argu- ments

PHP Command Script The **Command** value is run as PHP code

Restart Local Service The name of a local service to restart

Restart Remote Service The name of a remote service to restart via XMLRPC. This utilizes the system *XMLRPC sync configuration*

The GUI includes several examples of common actions

Certificate Renewal After When the package will attempt a renewal for the certificate. Default is `60` days (2 months). Certificates are valid for a maximum of `90` days.

Validation Methods

ACME providers can validate by checking the contents of a TXT record in DNS, or by fetching a file in a known location from a web server.

The ACME package support validating directly with standalone methods or webroot, but those options are less secure than DNS-based options. The ACME package also supports numerous methods to update various DNS providers. Wildcard certificates can only be obtained through DNS-based methods (*Wildcard Certificates*)

Tip: DNS-based update methods are the best practice as they does not require external inbound access. They can be used for internal systems that do not allow or cannot receive Internet traffic.

The following list is only a portion of the validation methods supported by the package.

nsupdate

The nsupdate method uses RFC 2136 style DNS updates to populate a TXT record in DNS.

Before starting, an appropriate DNS key and settings must be in place in the DNS infrastructure for the domain to allow the host to update a TXT DNS record for _acme-challenge.<domain name>.

This method has the following options:

Server The IP address or hostname of the DNS server to which the client sends updates

Key Name The name of the update key

Leave this blank unless it is different than _acme-challenge.<domain name> **Key Algorithm** The algorithm used for the key, which must match the key and the server **Key** The update key for this record

Zone Sets the zone name the package sends to the DNS server in the update request

DNS-Manual

The manual DNS method can be utilized when a firewall cannot receive inbound traffic and it does not have access to any automatic DNS-based method.

The **manual** in the name indicates that the process **must be performed by hand** both initially and when it is time to renew the certificate. The firewall obtains the authorization value and then the TXT record must be manually created or updated with this value.

Warning: Avoid using this method unless no other method is available.

To use this method:

- Add an entry to the **Domain SAN list**
- **Mode**: Enabled
- Enter domain name (e.g. myhost.example.com)
- Set **Method** to *DNS-Manual*
- Click **Save**
- Click **Issue**
- Locate the record info in the output:

```
[Mon Feb 6 14:49:23 EST 2017] Add the following TXT record:
[Mon Feb 6 14:49:23 EST 2017] Domain: '_acme-challenge.www.example.com'
[Mon Feb 6 14:49:23 EST 2017] TXT value: 'xPrykHSri5epT5yrJJWyY536Z1T51r_
```

(continues on next page)

(continued from previous page)

```
[Mon Feb 6 14:49:23 EST 2017] Please be aware that you prepend _acme-challenge.
↪before your domain
[Mon Feb 6 14:49:23 EST 2017] so the resulting subdomain will be: _acme-challenge.
↪www.example.com
```

- Add or update the TXT record in the domain's DNS server for `_acme-challenge.<domain name>` with the TXT value from the output
- Wait approximately 2 minutes, or longer, for DNS to propagate
- Click **Renew**

Namecheap API

For certain accounts with Namecheap, API access may be obtained that allows remote manipulation of DNS records. This can be used with the ACME package to validate certificates for domains with DNS hosted at Namecheap using their BasicDNS servers.

> **Warning:** The Namecheap DNS API requires that the client read all records and then write them all back when making any change. This is potentially dangerous. Take a backup of all DNS records on the domain before attempting to use the API.

The first step is to request API access:

- Login to a Namecheap account
- Navigate to **Profile** > **Tools** under the account
- Look for **Namecheap API Access** under **Business & Dev Tools**
- If the status does not say **On**, then click **Manage** and change the slider to **On**.

> **Note:** API access must be approved by Namecheap. There are qualifications to meet, such as a specific number of domains or a balance on the account. Check the Namecheap API documentation for more information. The process is documented as taking 2 days, but may take longer. If API access is not enabled after several days, contact Namecheap support.

Once the API is enabled, then perform the following steps:

- Login to a Namecheap account
- Navigate to **Profile** > **Tools** under the account
- Look for **Namecheap API Access** under **Business & Dev Tools**
- Click **Manage**
- Note the API key for use in the ACME package
- Click **Edit** and add whitelisted IP addresses that can contact the API using this API key.

- Now setup the account in the ACME package:
- Add an entry to the **Domain SAN list**

- **Mode**: Enabled
- Enter domain name (e.g. `myhost.example.com`)
- Set **Method** to *DNS-Namecheap*
- Click + to expand the method-specific settings
- Fill in the info

  API Key The API Key displayed in the Namecheap API Access manager, as described previously.

  Username The Namecheap account username associated with the API Key.

- Ensure the other options are set properly, per *Create a certificate*.
- Click **Save**
- Click **Issue/Renew**

Other DNS Methods

The package contains several additional DNS-based methods for other providers. These work similar to the nsupdate method above, but have configuration values specific to each provider. Contact the DNS provider or server adminis- trator to obtain the necessary settings or credentials.

FTP Webroot

The **FTP webroot** method is useful when the firewall is performing NAT (port forward or 1:1) or reverse proxy duty for handling traffic for the domain. The firewall can use SFTP or FTPS to store the domain validation files on a web server behind the firewall so it does not have to host the files itself.

We recommend using this method when no DNS update method is available for use

by the firewall. This method has the following options:

Server The server where the package will send the challenge response files, e.g. `sftp://x.x.x.x`

Note: This method supports supports `sftp://` and `ftps://` servers.

Username/password Credentials for the SFTP/FTPS account

Folder Full path to the target directory including `/.well-known/acme-challenge` at the end

Warning: Make sure the specified user has write permissions to the directory!

Webroot Local Folder

This method works similar to FTP Webroot but with the files hosted on the firewall itself. This method cannot be utilized by the WebGUI web server as that would mean exposing the GUI to the Internet, which is a major security issue.

This method can, however, be used in conjunction with the HAProxy package to host the files on the firewall itself in some circumstances.

Standalone (HTTP/TLS-ALPN)

The **Standalone** methods for HTTP and TLS-ALPN run a small web server natively that is active only while the validation process is running. The TLS-ALPN method is more secure as it encrypts communication with the ACME provider.

Warning: We do not recommend using these methods as they expose a service on the firewall to the Internet. Only use these methods if no other method is available.

Warning: The service **must** be accessible using port `80` (HTTP) or `443` (TLS-ALPN)!

If the firewall is using port `80` (HTTP) or `443` (TLS-ALPN) for another service, such as the firewall GUI or its redirect, then this method may not be viable. If the service on the port is public, then it cannot be used. If the service is private, then it may be possible to relocate the existing service or bind the update method to an alternate port, then port forward on the WAN interface.

A firewall rule must allow traffic to the target port at all times, it cannot be automatically enabled and disabled in the current package.

Note: The standalone binding should only be changed if the port is forwarded via NAT to a different port (e.g. 80 forwarded to 8080)

Standalone HTTP

Standalone HTTP Server has the following options:

Port Port to which the package will bind listening for HTTP requests with a stand-alone server. Must be `80` or port `80` must be forwarded to this port on the default gateway WAN.

> **Warning:** If port `80` is used by the standalone service, the GUI redirect must be disabled on **System > Advanced** using the **Disable webConfigurator redirect rule** option. If the redirect is active when standalone mode attempts to use the port, it will print an error message stating that `socat` is unable to bind to the port.

Bind to IPv6 instead of IPv4 If the domain name for the firewall has both an A and AAAA DNS record, check this option so that validation can occur over IPv6.

Standalone TLS-ALPN

Standalone TLS-ALPN Server has the following options:

Port Port to which the package will bind listening for TLS-ALPN requests with a stand-alone server. Must be `443` or port `443` must be forwarded to this port on the default gateway WAN.

> **Warning:** If port `443` is used by the standalone service, the GUI must be moved to an alternate port on **System > Advanced** using the **TCP Port** option for the GUI. If the redirect is active when standalone mode attempts to use the port, the

DNS Alias Mode

DNS Alias mode allows a DNS update method to update an alternate domain name instead of updating a record for the domain name directly.

If the main DNS provider does not support updating TXT records, a CNAME record can point to an alternative domain which does.

Challenge Alias

In Challenge Alias mode (default), the ACME package still automatically prepends `_acme-challenge.` to both the **Domain Name** and the **DNS Alias** domain.

In the certificate entry, set:

Domain Name `company.example` which does not support automatic updates

DNS Alias Domain `dynamic.example` which is the alternative domain in a dynamic zone

DNS Domain Alias mode Leave unchecked

On the DNS server, add a CNAME record pointing to the **DNS Alias** hostname with `_acme-challenge.`
prepended:

```
_acme-challenge.company.example        IN      CNAME   _acme-challenge.dynamic.example.
```

When updating, the package will update `_acme-challenge.dynamic.example` in DNS while sending
`company.example` in the certificate request to the ACME provider.

Domain Alias

Domain Alias mode works similar to Challenge Alias mode but it **does not** prepend `_acme-challenge.` to the **DNS Alias** domain. Some administrators prefer this when using many hostnames in a single dynamic zone, or for working around limitations in DNS providers or platforms.

In the certificate entry, set:

Domain Name `company.example` which does not support automatic updates

DNS Alias Domain `checkme.dynamic.example` which is the alternative domain in a dynamic zone

DNS Domain Alias mode Checked

On the DNS server, add a CNAME record pointing directly to the **DNS Alias** hostname:

```
_acme-challenge.company.example        IN      CNAME   checkme.dynamic.example.
```

When updating, the package will update `checkme.dynamic.example` in DNS while sending `company. example` in the certificate request to the ACME provider.

General Settings

These settings control the general behavior of the ACME package and are not specific to any single certificate or key.

Cron Entry A checkbox which enables the ACME renewal cron job. When set, the ACME package will check all certificates each night and if any are up for renewal, it will attempt to renew them.

Write Certificates When set, the ACME package will write the certificate files out in `/conf/acme`. From there, other scripts or processes which do not support GUI integration can pick up the certifi- cate.

24.1.4Wildcard Certificates

Let’s Encrypt supports wildcard certificates (e.g. *.example.com) with their ACMEv2 infrastructure. A wildcard certificate will work for any hostname inside a given domain, which helps with handling certificates for multiple domains.

Note: Unrelated to ACME, but wildcard certificates in general: A wildcard only helps for **one level** of subdo- mains. For example, *.example.com will work for host.example.com but will NOT work for host.sub. example.com. If hosts are structured in this way, a wildcard certificate is required for each sub zone, e.g. *.sub. example.com.

Wildcard validation **requires a DNS-based method** and works similar to validating a regular domain. For example, to get a certificate for *.example.com, the package updates a TXT record in DNS the same as it would for example. com, which means the DNS record (and potentially key name) would be for _acme-challenge.example.com.

To obtain a wildcard certificate, follow the same procedures as other DNS validation methods, with the following differences:

- The **Account Key** must be registered with an ACME v2 server (staging for testing, or production)
- The **Domain SAN list** should contain entries for the base domain (e.g. example.com and the wildcard version of the same domain (e.g. *.example.com. The settings will be the same for both entries.
- For *DNS-NSupdate / RFC 2136*: Set the **Key Name** to the base domain (example.com) for both entries.

24.2Cache / Proxy

Proxies are intermediaries that sit between clients and servers. A client connects to a proxy, and then the proxy decides if the client can receive content from a server. If so, the proxy makes its own connection to the server and then passes back data to the client.

There are two major types of proxies:

Forward Proxy Typically sits between local clients and remote Internet servers. It can be used to control which web sites that clients are allowed to load, or log servers and URLs clients are visiting. These mostly work with HTTP, but in special cases can also work with HTTPS.

Reverse Proxy Typically sits between remote clients and local servers. These allow for load balancing, failover, or other intelligent connection routing for public services such as web servers.

24.2.1Squid

Squid is primarily a forward proxy used for client access control. It can, however, be used in a reverse proxy role if needed. The reverse proxy capabilities are inferior to HAProxy, however.

The most common use case for squid is covered in *Configuring the Squid Package as a Transparent HTTP Proxy*. Additional documentation below covers related topics.

Configuring the SquidGuard Package

squidGuard is a URL redirector used to integrate blacklists with the Squid proxy software. There are two big advan- tages to squidGuard: it is fast and it is free. squidGuard is published under the GNU Public License.

squidGuard can be used to:

- Limit the web access for some users to a list of accepted/well known web servers and/or URLs only.
- Block access to some listed or blacklisted web servers and/or URLs for some users.
- Block access to URLs matching a list of regular expressions or words for some users.
- Enforce the use of domain names/prohibit the use of IP addresses in URLs.
- Redirect blocked URLs to an info page.
- Redirect banners to an empty GIF.
- Have different access rules based on time of day, day of the week, date etc.

Installing Squid and squidGuard

1. From the AZTCO-FW® webGUI, navigate to **System > Packages**, **Available Packages** tab
2. Install the **Squid** package if it is not already installed.
3. Install the **squidGuard** package
4. Configure **Squid** package.
5. Configure **squidGuard** package.

Configure the squidGuard Package

Basic configuration

Here describes how to enable and configure squidGuard, and common users access.

1. Open **General settings** tab.
    1. Check the **Enable** box to activate the package.
    2. Set **Blacklist** options to use blacklist categories. (See above, optional)
    3. Click **Save** button.
2. Open **Common ACL** page.
    1. Click **Target Rules List** to show defined blacklists and target categories
        1. Define default user access: select **Default access [all]** as *allow* or *deny*.
        2. Define other category actions:
            1. Select **—**, to ignore a category.
            2. Select **allow**, to allow this category for clients.

   3. Select **deny**, to deny this category for clients.

   4. Select **white**, to allow this category without any restrictions. This option is used for exceptions to prohibited categories.

  3. To prohibit clients from using IP addresses in URLs, check **Do Not Allow IP Addresses in URL**.

  4. Select **Redirect mode**:

     1. *Int error page*: Use the built-in error page. A custom message may be entered in the **Redirect info** box below.

     2. *Int blank page*: Redirect to a blank page

     3. The other options are various redirects to external error pages, and a URL must be entered in the **Redirect info** box if they are chosen.

  5. **Use safe search engine**: Protect customers from unwanted search results. It is supported by *Google, Yandex, Yahoo, MSN, Live Search*. Make sure that these search engines are available. If this protection should be strictly enforced, disable access to all other search engines.

3. After settings are complete, return to the **General Settings** tab and press **Apply**.

Blacklist

Blacklists are optional, but often useful for allowing access to certain types of sites.

squidGuard comes with a small blacklist basically for testing purposes. They should not be used in production. A better way is to start with one of the blacklist collections listed (alphabetically) below.

- MESD blacklists - They are freely available.
- Shalla's Blacklists - Free for non commercial/private use. (Recommended)
- more..

Downloading
blacklists:

1. Open **General Settings** tab in squidGuard package GUI, found at **Services > Proxy Filter**.
2. Check **Blacklist** to enable the use of blacklists.
3. Enter blacklist URL in the field **Blacklist URL**.
4. If the firewall is itself behind a proxy, enter the proxy information in **Blacklist proxy** (this step is not necessary for most people).
5. Click **Save**.
6. Navigate to the **Blacklist** tab inside of squidGuard.
7. Click the **Download** button.
8. Wait while blacklist will downloaded and prepared to use (10-35 min). Progress will be displayed on that page as the list is downloaded and processed.

How-Tos

Exclude domain/URL from blacklist

In the squidGuard GUI (**Services > Proxy Filter**):

1. Open the **Target categories** page
2. Click ➕ to add a new item
3. Enter a name for the category - `myWhitelist` for example.
4. Add domains and/or URLs to the lists as needed. Entries should be separated by a space. The examples on the page show how entries should be formatted.
5. As with the Common ACL discussed previously, redirect and logging options specific to this category may be set.
6. Click **Save**.
7. Open **Common ACL** or **Groups ACL** page (whichever should have an exclusion).
8. Click **Target Rule List** to expand the list of categories. The newly created category should show alphabetically in the list, above any blacklist categories. Find the **MyWhiteList** entry in the list and select **white**.
9. Click **Save**.
10. Return to the **General Settings** tab and press **Apply**.

Block download by Extension

In the squidGuard GUI (**Services > Proxy Filter**):

1. Open the **Target categories** page.
2. Click ➕ to add a new item.
3. Enter a name for the category - `myBlockExt` for example.
4. Add Expressions (for example for asf, zip, exe and etc files):

```
(.*\/.*\.(asf|wm|wma|wmv|zip|rar|cab|mp3|avi|mpg|swf|exe|mpeg|mp.|mpv|mp3|wm.
→|vpu))
```

5. Click **Save**.
6. Open **Common ACL** or **Groups ACL** page (whichever should have an exclusion).
7. Click **Target Rule List** to expand the list of categories. The newly created category should show alphabetically in the list, above any blacklist categories. Find the **myBlockExt** entry in the list and select **deny**.
8. Click **Save**.
9. Return to the **General Settings** tab and press **Apply**.

10.

24.3FreeRADIUS package

FreeRADIUS is a free implementation of the RADIUS protocol. Supports MySQL, PostgreSQL, LDAP, Kerberos. Refer to the following articles for more information on the listed topics:

24.3.1Features

The features below were tested on AZTCO-FW software version 2.x

- Authentication with Captive-Portal
- Pre-defined user attributes and custom check-items and reply-items
- NAS/Clients running on IPv4 and IPv6
- Interfaces can listen on IPv4 and IPv6
- OpenVPN + Username + RADIUS and OpenVPN + Username + Cert + RADIUS
- Auth with PAP, CHAP, MSCHAP, MSCHAPv2
- Auth with EAP-MD5 + dynamic VLAN assignment
- Auth with PEAP + dynamic VLAN assignment
- Auth with EAP-TLS/EAP-TTLS + dynamic VLAN assignment

```
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense_
→port 0 cli 00-04-23-5C-9D-19)
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense
→port 0 cli 00-04-23-5C-9D-19)
radiusd[3206]: Login OK: [testuser/<via Auth-Type = EAP>] (from client pfsense
→port 0 via TLS tunnel)
```

- Simultaneous-Use - The following will be present in the system log

```
radiusd[3206]: Multiple logins (max 1) : [testuser/testpw] (from client testing_
→port 10)
```

- A certain amount of time per day/week/month/forever (`CHECK-ITEM: Max-Daily-Session := 60`) The user will be disconnected and cannot re-login after the amount of time is reached:

```
radiusd[3206]: Invalid user (rlm_counter: Maximum daily usage time reached):_
→[testuser/<via Auth-Type = EAP>] (from client pfsense port 0 cli 00-04-23-5C-9D-
```

- A certain amount of traffic per day/week/month/forever. The user will be disconnected and cannot re-login after the amount of traffic is reached. The syslog output looks like this:

```
root: FreeRADIUS: Used amount of daily upload and download traffic by testuser is_
→0 of 100 MB! The user was accepted!!!
root: FreeRADIUS: Credentials are probably correct but the user testuser has
→reached the daily amount of upload and download traffic which is 243 of 100 MB!
```

- MySQL

-
- LDAP/ActiveDirectory (connecting to MS AD with PAP)
- User-Auth with SQUID
- One-Time-Password

24.3.2Installation and Configuration

- Navigate to **System** > **Packages**, **Available Packages** tab.
- Click + at the end of the row for**freeradius3**.
- Confirm the installation.
- Monitor the progress as it installs.

After Installation, the service may be configured at **Services > FreeRADIUS**.

- Select the **interface(s)** on which the RADIUS server should listen on.
- Configure the **NAS/client(s)** from which the RADIUS server should accept packets.
- Add the **user(s)** who should have access.

After this, have a look at the AZTCO-FW® syslog. There should be the following:

```
radiusd[16634]: Ready to process requests.
radiusd[16627]: Loaded virtual server
```

Troubleshooting RADIUS Authentication

When attempting to authenticate against a RADIUS server, errors may be encountered in the logs that prevent it from working properly. Here are some errors and how to resolve them:

```
mpd: [pt0] RADIUS: RadiusSendRequest: rad_init_send_request failed: -1
```

- This appears to happen when the RADIUS shared secret contains special characters. Try again with an alphanumeric shared secret.

Get FreeRADIUS Status Server Updates

The status server will give lots of information about the FreeRADIUS server. Many stats are shown about Accounting- Packets, dropped packets and much more. To enable status server and request information from the server do the following:

- Setup an interface with **Interface-Type**: *status* and a free port. The default port for RADIUS accounting is `1813`.
- Setup a NAS/Client with **IP-Address**: `127.0.0.1` and a password. Password *testing123* will be used in this example.
- SSH to the AZTCO-FW firewall and enter the following command on the command line:

```
echo "Message-Authenticator = 0x00, FreeRADIUS-Statistics-Type = All" | \
→radclient localhost:1813 status testing123
```

The output should look like this:

```
Received response ID 223, code 3, length = 140
       FreeRADIUS-Total-Access-Requests = 1
       FreeRADIUS-Total-Access-Accepts = 0
       FreeRADIUS-Total-Access-Rejects = 14
       FreeRADIUS-Total-Access-Challenges = 0
       FreeRADIUS-Total-Auth-Responses = 14
       FreeRADIUS-Total-Auth-Duplicate-Requests = 0
```

To request other status updates, replace **FreeRADIUS-Statistics-Type = 1** from the command above with another value. More values can be found in this path on the AZTCO-FW firewall:

```
/usr/local/share/freeradius/dictionary.freeradius
```

24.4OpenVPN Client Export Package

The easiest way to configure an OpenVPN client on most platforms is to use the OpenVPN Client Export Package on the AZTCO-FW® firewall.

Install the OpenVPN Client Export Utility package as follows:

- Navigate to **System** > **Packages**
- Locate the **OpenVPN Client Export** package in the list
- Click + **Install** next to that package listing to install

Once installed, it can be found at **VPN > OpenVPN**, on the **Client Export** tab.

The options for the package include:

Remote Access Server Pick the OpenVPN server instance for which a client will be exported. If there is only one OpenVPN remote access server there will only be one choice in the list. The list will be empty if there are no Remote Access mode OpenVPN servers.

Host Name Resolution Controls how the “remote” entry the client is formatted.

Interface IP Address When chosen, the interface IP address is used directly. This is typ- ically the best choice for installations with a static IP address on WAN.

Automagic Multi-WAN IPs This option is useful when redirecting multiple ports using port forwards for deployments that utilize multi-WAN or multiple ports on the same WAN. It will seek out and make entries for all port forwards that target the server and use the destination IP address used on the port forward in the client configuration.

Automagic Multi-WAN DDNS Hostnames Similar to the previous option, but it uses the first Dynamic DNS entry it finds that matches the chosen destination.

Installation Hostname Places the firewall's hostname, defined under **System > General Setup**, into the client configuration. The hostname must exist in public DNS so it can be resolved by clients.

Dynamic DNS Hostname Entries Each Dynamic DNS hostname configured on the fire- wall is listed here. These are typically the best choice for running a server on a single WAN with a dynamic IP address.

Other Presents a text box in which a hostname or IP address can be entered for the client to use.

Verify Server CN Specifies how the client will verify the identity of the server certificate. The CN of the server certificate is placed in the client configuration, so that if another valid certificate pretends to be the server with a different CN, it will not match and the client will refuse to connect.

Automatic - Use verify-x509-name where possible This is the best for current clients. Older methods have been deprecated since this method is more accurate and flexible.

Use tls-remote This can work on older clients (OpenVPN 2.2.x or earlier) but it will break newer clients as the option has been deprecated.

Use tls-remote and quote the server CN Works the same as *tls-remote* but adds quotes around the CN to help some clients cope with spaces in the CN.

Do not verify the server CN Disables client verification of the server certificate common name.

Use Random Local Port For current clients, the default (checked) is best, otherwise two OpenVPN con- nections cannot be run simultaneously on the client device. Some older clients do not support this, however.

Use Microsoft Certificate Storage Under Certificate Export Options, for exported installer clients this will place the CA and user certificate in Microsoft's certificate storage rather than using the files directly.

Use a password to protect the pkcs12 file contents When checked, enter a Password and confirm it, then the certificates and keys supplied to the client will be protected with a password. If the Open- VPN server is configured for user authentication this will cause users to see two different password prompts when loading the client: One to decrypt the keys and certificates, and another for the server's user authentication upon connecting.

Use Proxy If the client will be located behind a proxy, check *Use proxy to communicate with the server* and then supply a Proxy **Type**, **IP Address**, **Port**, and **Proxy Authentication** with credentials if needed.

OpenVPNManager When checked, this option will bundle the Windows installer with OpenVPNMan- ager GUI in addition to the normal Windows client. This alternate GUI manages the OpenVPN service in such a way that it does not require administrator-level privileges once installed.

Additional configuration options Any extra configuration options needed for the client may be placed in this entry box. This is roughly equivalent to the **Advanced options** box on the OpenVPN config- uration screens, but from the perspective of the client.

Note: There is no mechanism to save these settings, so they must be checked and set each time the page is visited.

24.4.1Client Install Packages List

Under **Client Install Packages** is a list of potential clients to export. The contents of the list depend on how the server is configured and which users and certificates are present on the firewall.

The following list describes how the server configuration style affects the list in the package:

Remote Access (SSL/TLS) User certificates are listed which are made from the same CA as the OpenVPN server

Remote Access (SSL/TLS + User Auth – Local Users) User entries are listed for local users which also have an associated certificate made from the same CA as the OpenVPN server.

Remote Access (SSL/TLS + User Auth – Remote Authentication) Because the users are remote, user certificates are listed which are made from the same CA as the OpenVPN server. It is assumed that the username is the same as the common name of the certificate.

Remote Access (User Auth – Local Users or Remote Authentication) A single configuration entry is shown for all users since there are no per-user certificates.

The example setup from the wizard made previously in this chapter was for SSL/TLS + User Auth with Local Users, so one entry is shown per user on the system which has a certificate created from the same CA as the OpenVPN server.

Note: If no users are shown, or if a specific user is missing from the list, the user does not exist or the user does not have an appropriate certificate. See *Local Users* for the correct procedure to create a user and certificate.

24.4.2Client Install Package Types

Numerous options are listed for each client that export the configuration and associated files in different ways. Each one accommodates a different potential client type.

Standard Configurations

Archive Downloads a ZIP archive containing the configuration file, the server's TLS key if defined, and a PKCS#12 file which contains the CA certificate, client key, and client certificate. This option is usable with Linux clients or Tunnelblick, among others.

File Only Downloads only the basic configuration file, no certificates or keys. This would mainly be used to see the configuration file itself without downloading the other information.

24.5pfBlocker-NG Package

pfBlocker-NG introduces an Enhanced Alias Table Feature to

AZTCO-FW® software. What it allows:

- Assigning many IP address URL lists from sites like I-blocklist to a single alias and then choose a rule action.
- Blocking countries and IP ranges.
- Replacement of both **Countryblock** and **IPblocklist** by providing the same functionality, and more, in one package.
- Uses native functions of AZTCO-FW software instead of file hacks and table

manipulation. Features include:

- Country_Block features
- IP_Blocklist features
- Dashboard widget
- XMLRPC Sync
- Dashboard widget with aliases applied and package hit
- Lists update frequency
- Many new options to choose what to block and how to block.
- Network lists may be used for custom rules.

24.5.1General Setup

Set the interfaces to be monitored by pfBlocker-NG (both inbound and outbound), where the inbound is the Internet connection.

To prevent devices or users from accessing sites in the selected countries/IP addresses, select local interfaces under
outbound.

29.17.1 Setting up Lists

This is the IPBlocklist feature, enter IP addresses here to specifically block. It must be in the file format or CIDR. Create a list for each type of action to be taken by pfBlocker.

Options are:

`Deny Both` - Will deny access on Both directions.

`Deny Inbound` - Will deny access from selected lists to the local network.

`Deny Outbound` - Will deny access from local users to IP address lists selected to block.

`Permit Inbound` - Will allow access from selected lists to the local network.

`Permit Outbound` - Will allow access from local users to IP address lists selected to block.

`Disabled` - Will just keep selection and do nothing to selected Lists.

`Alias Only` - Will create an alias with selected Lists to help custom rule assignments.

The rest of the tabs (except sync) specify the other lists included with the package. They are separated by continent with the exception of the spammer list which contains countries from around the globe that are known to harbor spammers.

Sync tab configures pfBlocker to sync its configuration to other AZTCO-FW devices.

24.6IDS / IPS

AZTCO-FW® software can act in an Intrusion Detection System (IDS) / Intrusion Prevention System (IPS) role with add-on packages like Snort and Suricata.

Note: The Snort and Suricata packages share many design similarities, so in most cases the instructions for Snort carry over to Suricata with only minor adjustments.

24.6.1Snort

Configuring the Snort Package

Snort is an intrusion detection and prevention system. It can be configured to simply log detected network events to both log and block them. Thanks to OpenAppID detectors and rules, Snort package enables application detection and filtering. Snort operates using detection signatures called rules. Snort rules can be custom created by the user, or any of several pre-packaged rule sets can be enabled and downloaded.

The Snort package currently offers support for these pre-packaged rules:

- AZTCO-FW Snort rules
- AZTCO FW Emerging Threats Open Rules
- OpenAppID Open detectors and rules for application detection

Launching Snort configuration GUI

To launch the Snort configuration application, navigate to **Services > Snort** from the menu in the AZTCO-FW webGUI.

Setting up Snort package for the first time

Click the **Global Settings** tab and enable the rule set downloads to use

Once the desired rule sets are enabled, next set the interval for Snort to check for updates to the enabled rule packages. Use the **Update Interval** drop-down selector to choose a rule update interval. In most cases every *12* hours is a good choice. The update start time may be customized if desired. Enter the time as hours and minutes in 24-hour time format. The default start time is *3* minutes past midnight local time. So with a *12*-hour update interval selected, Snort will check the

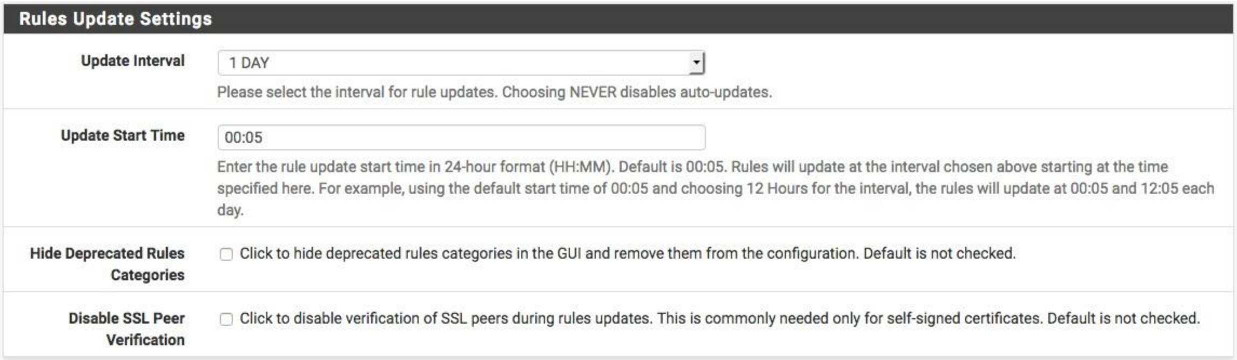

Rules Update Settings

Update Interval: 1 DAY
Please select the interval for rule updates. Choosing NEVER disables auto-updates.

Update Start Time: 00:05
Enter the rule update start time in 24-hour format (HH:MM). Default is 00:05. Rules will update at the interval chosen above starting at the time specified here. For example, using the default start time of 00:05 and choosing 12 Hours for the interval, the rules will update at 00:05 and 12:05 each day.

Hide Deprecated Rules Categories: □ Click to hide deprecated rules categories in the GUI and remove them from the configuration. Default is not checked.

Disable SSL Peer Verification: □ Click to disable verification of SSL peers during rules updates. This is commonly needed only for self-signed certificates. Default is not checked.

Update the rules

The **Updates** tab is used to check the status of downloaded rules packages and to download new updates. The table shows the available rule packages and their current status (not enabled, not downloaded, or a valid MD5 checksum and date).

Click on the **Update Rules** button to download the latest rule package updates. If there is a newer set of packaged rules on the vendor web site, it will be downloaded and installed. The determination is made by comparing the MD5 of the local file with that of the remote file on the vendor web site. If there is a mismatch, a new file is downloaded. The **FORCE** button can be used to force download of the rule packages from the vendor web site no matter how the MD5 hash tests out.

Add Snort to an interface

Click the **Snort Interfaces** tab and then the ➕ icon to add a new Snort interface.

A new Interface Settings tab will open with the next available interface automatically selected. The interface selection may be changed using the **Interface** drop-down if desired. A descriptive name may also be provided for the interface. Other interface parameters may also be set on this page. Be sure to click the **SAVE** button down at the bottom of the page when finished.

After saving, the browser will be returned to the **Snort Interfaces** tab. Note the warning icons in the image below showing no rules have been selected for the new Snort interface. Those rules will be configured next. Click the ✏ icon (shown highlighted with a red box in the image below) to edit the new Snort interface again.

Select which types of rules will protect the network

Click the **Categories** tab for the new interface.

Be sure to click **SAVE** when finished to save the selection and build the rules file for Snort to use.

Starting Snort on an interface

Click the **Snort Interfaces** tab to display the configured Snort interfaces. Click the ⊗ icon (shown highlighted with a red box in the image below) to start Snort on an interface.

It will take several seconds for Snort to start. Once it has started, the icon will change to ✔ as To stop a running Snort instance on an interface, click the ✔ icon.

Select which types of signatures will protect the network

Click the **Rules** tab for the interface to configure individual rules in the enabled categories. Generally this page is only used to disable particular rules that may be generating too many false positives in a particular network environment. Be sure they are in fact truly false positives before taking the step of disabling a Snort rule!

Select a rules category from the **Category** drop-down to view all the assigned rules.

Define servers to protect and improve performance

Managing blocked hosts

The **Blocked** tab shows what hosts are currently being blocked by Snort (when the block offenders option is selected on the **Interface Settings** tab). Blocked hosts can be automatically cleared by Snort at one of several pre-defined intervals. The blocking options for an interface are configured on the Snort **Interface Settings** tab for the interface.

Managing Pass lists

Pass Lists are lists of IP addresses that Snort should never block. These may be created and managed on the **Pass Lists** tab. When an IP address is listed on a Pass List, Snort will never insert a block on that address even when malicious traffic is detected.

To create a new Pass List, click ➕. To edit an existing Pass List, click the ✏. To delete a Pass click 🗑

. Note that a Pass List may not be deleted if it is currently assigned to one or more Snort interfaces.

A default Pass List is automatically generated by Snort for every interface, and this default list is used when no other list is specified. Pass Lists are assigned to an interface on the **Interface Settings** tab.

Customized Pass List may be created and assigned to an interface. This might be done when trusted external hosts exist that are not located on networks directly connected to the firewall. To add external hosts in this manner, first create an Alias under **Firewall** > **Aliases** and then assign that alias to the **Assigned Aliases** field. In the example shown below, the alias "*Friendly_ext_hosts*" has been assigned. This alias would contain the IP addresses of the trusted external hosts.

When creating a custom Pass List, leave all the auto-generated IP addresses checked in the **Add auto-generated IP**

addresses section. Not selecting the checkboxes in this section can lead to blocking of critical addresses including the firewall interfaces themselves. This could result in being locked out of the firewall over the network! Only uncheck boxes in this section when absolutely necessary.

Click the **ALIASES** button to open a window showing previously defined aliases for selection. Remember to click
SAVE to save changes.

Note: Remember that simply creating a Pass List is only the first step! It must be selected by going to the **Interface Settings** tab for the Snort interface and assigning the newly created Pass List as shown below. After assigning and saving the new Pass List, restart Snort on the affected interface to pick up the change.

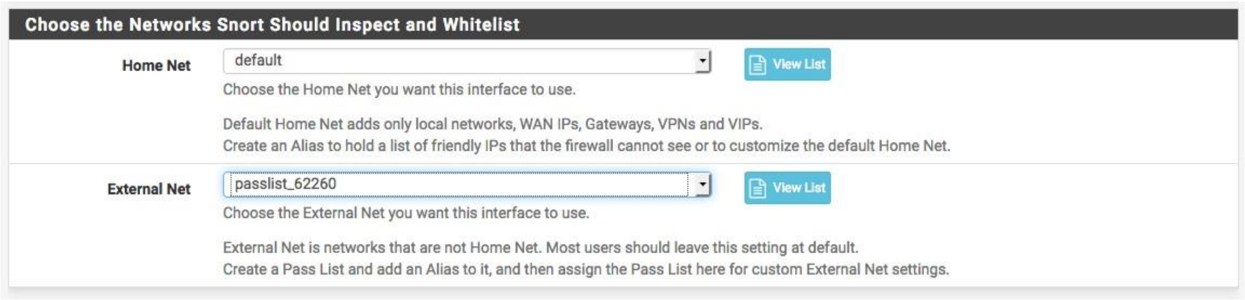

Choose the Networks Snort Should Inspect and Whitelist

Home Net: default

Choose the Home Net you want this interface to use.

Default Home Net adds only local networks, WAN IPs, Gateways, VPNs and VIPs.
Create an Alias to hold a list of friendly IPs that the firewall cannot see or to customize the default Home Net.

External Net: passlist_62260

Choose the External Net you want this interface to use.

External Net is networks that are not Home Net. Most users should leave this setting at default.
Create a Pass List and add an Alias to it, and then assign the Pass List here for custom External Net settings.

Alert Thresholding and Suppression

Suppression Lists allow control over the alerts generated by Snort rules. When an alert is suppressed, then Snort no longer logs an alert entry (or blocks the IP address if block offenders is enabled) when a particular rule fires. Snort still inspects all network traffic against the rule, but even when traffic matches the rule signature, no alert will be generated. This is different from disabling a rule. When a rule is disabled, Snort no longer tries to match it to any network traffic. Suppressing a rule might be done in lieu of disabling the rule when alerts should only be stopped based on either the source or destination IP. For example, to suppress the alert when traffic from a particular trusted IP address is the source. If any other IP is the source or destination of the traffic, the rule would still fire. To eliminate all alerts from the rule, then it is more efficient to simply disable the rule rather than to suppress it. Disabling the rule will remove it from Snort's list of match rules and therefore makes for less work Snort has to do.

On the Suppress List Edit page, a new suppress list entry may be manually added or edited. It is usually easier and faster to add suppress list entries by clicking ➕ shown with the alert entries on the **Alerts** tab. Remember to click the **SAVE** button to save changes when manually editing Suppress List entries.

Getting to know the alerts

The **Alerts** tab is where alerts generated by Snort are viewed. If Snort is running on more than one interface, choose the interface whose alerts should be viewed in the drop-down selector.

Use the **DOWNLOAD** button to download a gzip tar file containing all of the logged alerts to a local machine. The
CLEAR button is used to erase the current alerts log.

Alert Details

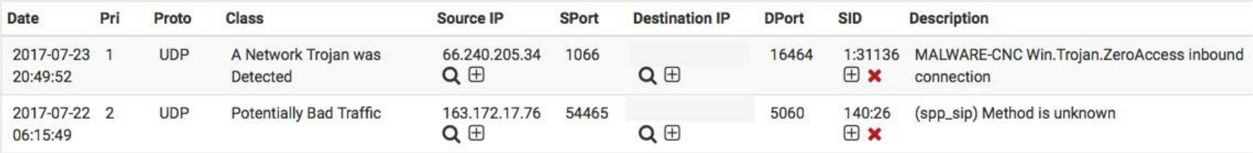

Date	Pri	Proto	Class	Source IP	SPort	Destination IP	DPort	SID	Description
2017-07-23 20:49:52	1	UDP	A Network Trojan was Detected	66.240.205.34 🔍 ⊞	1066	🔍 ⊞	16464	1:31136 ⊞ ✖	MALWARE-CNC Win.Trojan.ZeroAccess inbound connection
2017-07-22 06:15:49	2	UDP	Potentially Bad Traffic	163.172.17.76 🔍 ⊞	54465	🔍 ⊞	5060	140:26 ⊞ ✖	(spp_sip) Method is unknown

The **Date** column shows the date and time the alert was generated. The remaining columns show data from the rule that generated the alert.

In the **Source**, **Destination** columns are 🔍 icons for performing reverse DNS lookups on the IP addresses as well as a ➕ icon used to add an automatic *Suppress List* entry for the alert using the IP address and SID (signature ID). This will prevent future alerts from being generated by the rule for that specific IP address only. If either of the Source or Destination addresses are currently being blocked by Snort, then a 🗑 icon will also be shown. Clicking that icon will remove the block for the IP address.

The SID column contains two icons. The ✚ icon will automatically add that SID to the *Suppress List* for the interface and suppress future alerts from the signature for all IP addresses.

The ✖ icon in the SID column will disable the rule and remove it from the enforcing rule set.

When a rule is manually disabled, the icon in the SID column changes to ⊗.

Application ID detection with OpenApp ID

OpenAppID is an application-layer network security plugin for the open source intrusion detection system Snort. Learn more about it here.

Enabling OpenAppID and its rules is done from Snort **Global Settings**. Select both checkboxes to enable detectors and rules download. Save the page.

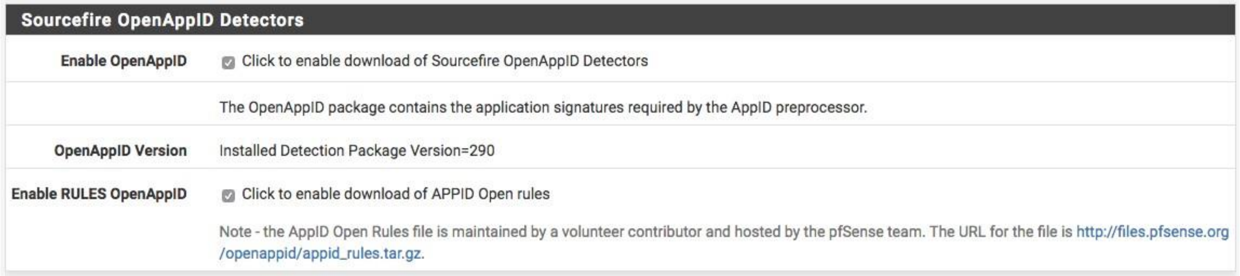

Sourcefire OpenAppID Detectors

Enable OpenAppID	☑ Click to enable download of Sourcefire OpenAppID Detectors
	The OpenAppID package contains the application signatures required by the AppID preprocessor.
OpenAppID Version	Installed Detection Package Version=290
Enable RULES OpenAppID	☑ Click to enable download of APPID Open rules Note - the AppID Open Rules file is maintained by a volunteer contributor and hosted by the pfSense team. The URL for the file is http://files.pfsense.org/openappid/appid_rules.tar.gz.

After enabling the detectors and rules go to Snort Updates tab and click on **Update Rules**. Wait for all the rules to update. Once done, the page will show OpenAppID detectors and rules have been updated.

The following steps assume the firewall already has a Snort interface for LAN. Edit the LAN interface and navigate to LAN categories tab. When there, make sure the **Snort OPENAPPID Rules** from the right column are all selected and click **Save**.

Scroll down to **Application ID Detection** section and select both **Enable** and **AppID Stats Logging** checkboxes. Save the page the OpenApp ID will be activated on the Snort interface.

Viewing detected applications can be done from **Alerts** tab. The following screenshots are examples of identified services and applications:

Facebook

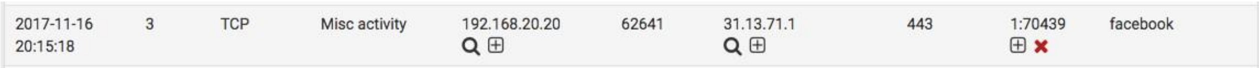

2017-11-16 20:15:18	3	TCP	Misc activity	192.168.20.20	62641	31.13.71.1	443	1:70439	facebook

Netflix

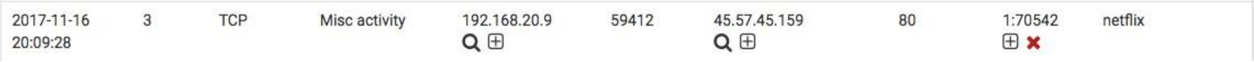

2017-11-16 20:09:28	3	TCP	Misc activity	192.168.20.9 🔍 ⊞	59412	45.57.45.159 🔍 ⊞	80	1:70542 ⊞ ✖	netflix

Reddit

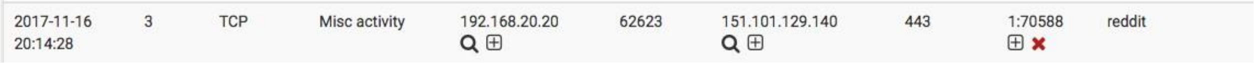

2017-11-16 20:14:28	3	TCP	Misc activity	192.168.20.20 🔍 ⊞	62623	151.101.129.140 🔍 ⊞	443	1:70588 ⊞ ✖	reddit

Amazon Web Services

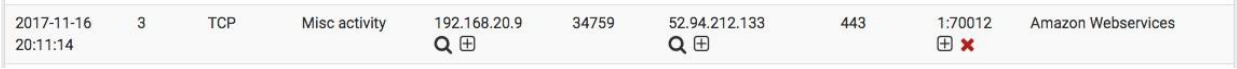

2017-11-16 20:11:14	3	TCP	Misc activity	192.168.20.9 🔍 ⊞	34759	52.94.212.133 🔍 ⊞	443	1:70012 ⊞ ✖	Amazon Webservices

iCloud

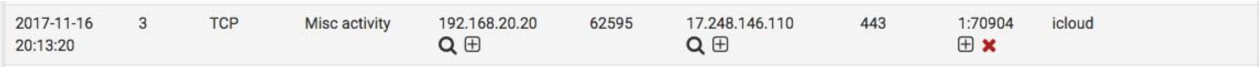

2017-11-16 20:13:20	3	TCP	Misc activity	192.168.20.20 🔍 ⊞	62595	17.248.146.110 🔍 ⊞	443	1:70904 ⊞ ✖	icloud

Twitter

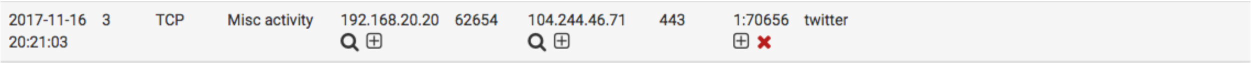

2017-11-16 20:21:03	3	TCP	Misc activity	192.168.20.20 🔍 ⊞	62654	104.244.46.71 🔍 ⊞	443	1:70656 ⊞ ✖	twitter

Snort Alerts

The **Alerts** tab is where alerts generated by Snort may be viewed. If Snort is running on more than one interface, choose the interface to view alerts for in the drop-down selector.

Use the **DOWNLOAD** button to download a gzip tar file containing all of the logged alerts to a local machine. The
CLEAR button is used to erase the current alerts log.

Alert Details

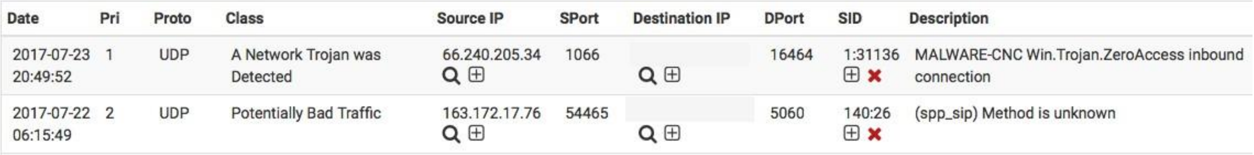

Date	Pri	Proto	Class	Source IP	SPort	Destination IP	DPort	SID	Description
2017-07-23 20:49:52	1	UDP	A Network Trojan was Detected	66.240.205.34 🔍 ⊞	1066	🔍 ⊞	16464	1:31136 ⊞ ✖	MALWARE-CNC Win.Trojan.ZeroAccess inbound connection
2017-07-22 06:15:49	2	UDP	Potentially Bad Traffic	163.172.17.76 🔍 ⊞	54465	🔍 ⊞	5060	140:26 ⊞ ✖	(spp_sip) Method is unknown

The **Date** column shows the date and time the alert was generated. The remaining columns show data from the rule that generated the alert.

In the **Source**, **Destination** columns are 🔍 icons for performing reverse DNS lookups on the IP addresses as well as a ⊞ icon used to add an automatic *Suppress List* entry for the alert using the IP address and SID (signature ID). This will prevent future alerts from being generated by the rule for that specific IP address only. If either of the Source or Destination addresses are currently being blocked by Snort, then a ✖ icon will also be shown. Clicking that icon will remove the block for the IP address.

The SID column contains two icons. The ⊞ icon will automatically add that SID to the *Suppress List* for the interface and suppress future alerts from the signature for all IP addresses.

The ✖ icon in the SID column will disable the rule and remove it from the enforcing rule set.

When a rule is manually disabled, the icon in the SID column changes to ⊗.

Snort Blocked Hosts

The **Blocked** tab shows what hosts are currently being blocked by Snort (when the **block offenders** option is selected on the **Interface Settings** tab). Blocked hosts can be automatically cleared by Snort at one of several pre-defined intervals. The blocking options for an interface are configured on the Snort **Interface Settings** tab for the interface. To manually remove a blocked host, click the 🗑 icon in the right-hand column.

The 🔍 icon performs a reverse DNS lookup on the blocked host IP address when clicked.

Snort Server Definitions

Define servers to protect and improve performance

The **Variables** tab is where the specific types of hosts on the network are configured. For example, the specific IP addresses or network ranges containing web servers to protect may be defined. This can make Snort more efficient because it won't waste time scanning for web server threats on IP addresses where web servers do not exist. Similarly, Snort performance can be optimized by instructing it which addresses contain other critical servers such as SMTP, POP, DNS, etc.

The exact ports or port ranges used for certain services on the network may also be specified.

Each value entered on this page can only be an existing Alias. Start typing the name of the Alias into a textbox and a drop-down selection of matching entries will appear for selection. Aliases are created under **Firewall > Aliases** from the menu.

Snort interface Settings

General Settings

Enable: used to enable or disable Snort on the selected interface. Snort is enabled on the interface when this box is checked.

Interface: used to choose which physical firewall interface this Snort instance protects.

Description: used to provide an optional friendly name for the interface.

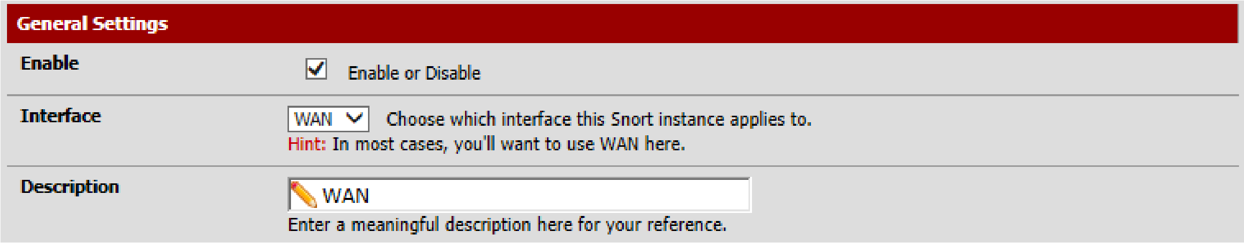

General Settings

Enable	☑ Enable or Disable
Interface	WAN Choose which interface this Snort instance applies to. Hint: In most cases, you'll want to use WAN here.
Description	WAN Enter a meaningful description here for your reference.

Alert Settings

Send Alerts to System Logs: when checked, all Snort alerts will be copied to the system log on the firewall.

Block Offenders: when checked, Snort will automatically insert a firewall block of the host generating an alert.

Kill States: when checked, Snort will kill all existing state table entries for the IP address it blocks. This should generally be enabled (box checked).

Which IP to Block: this determines which IP address extracted from the packet that generated an alert will be blocked. The choices are SOURCE, DESTINATION or BOTH. BOTH is the recommended default.

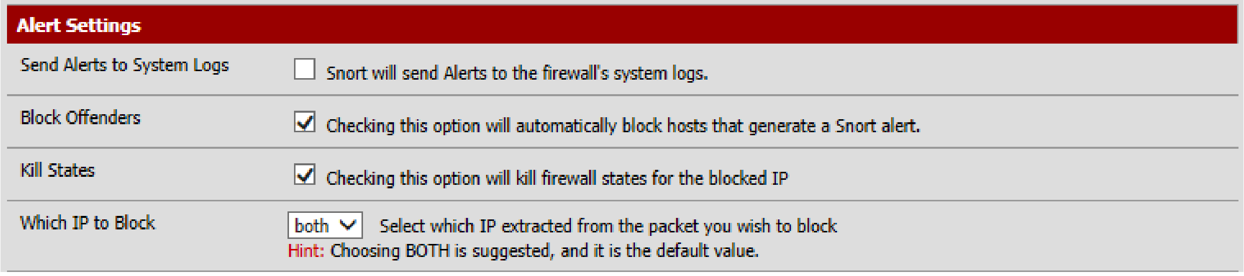

Alert Settings

Send Alerts to System Logs	□ Snort will send Alerts to the firewall's system logs.
Block Offenders	☑ Checking this option will automatically block hosts that generate a Snort alert.
Kill States	☑ Checking this option will kill firewall states for the blocked IP
Which IP to Block	both Select which IP extracted from the packet you wish to block Hint: Choosing BOTH is suggested, and it is the default value.

Detection Performance Settings

Search Method: used to select the pattern matcher algorithm used by Snort in the signature detection engine.

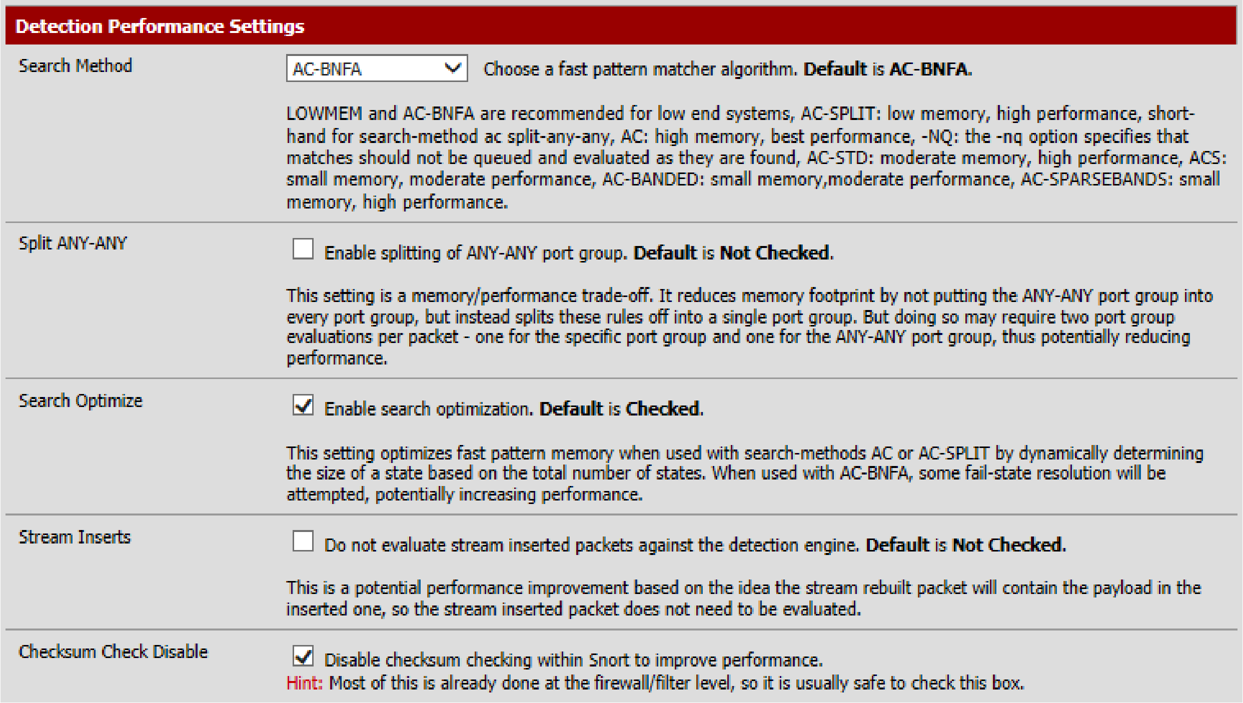

Detection Performance Settings

Search Method: AC-BNFA Choose a fast pattern matcher algorithm. **Default** is **AC-BNFA**.

LOWMEM and AC-BNFA are recommended for low end systems, AC-SPLIT: low memory, high performance, short-hand for search-method ac split-any-any, AC: high memory, best performance, -NQ: the -nq option specifies that matches should not be queued and evaluated as they are found, AC-STD: moderate memory, high performance, ACS: small memory, moderate performance, AC-BANDED: small memory,moderate performance, AC-SPARSEBANDS: small memory, high performance.

Split ANY-ANY: □ Enable splitting of ANY-ANY port group. **Default** is **Not Checked**.

This setting is a memory/performance trade-off. It reduces memory footprint by not putting the ANY-ANY port group into every port group, but instead splits these rules off into a single port group. But doing so may require two port group evaluations per packet - one for the specific port group and one for the ANY-ANY port group, thus potentially reducing performance.

Search Optimize: ☑ Enable search optimization. **Default** is **Checked**.

This setting optimizes fast pattern memory when used with search-methods AC or AC-SPLIT by dynamically determining the size of a state based on the total number of states. When used with AC-BNFA, some fail-state resolution will be attempted, potentially increasing performance.

Stream Inserts: □ Do not evaluate stream inserted packets against the detection engine. **Default** is **Not Checked**.

This is a potential performance improvement based on the idea the stream rebuilt packet will contain the payload in the inserted one, so the stream inserted packet does not need to be evaluated.

Checksum Check Disable: ☑ Disable checksum checking within Snort to improve performance.
Hint: Most of this is already done at the firewall/filter level, so it is usually safe to check this box.

Choose the networks Snort should inspect and whitelist

Home Net: selects the network Snort will use as the HOME_NET variable. Default is the recommended choice and contains the firewall WAN IP address and WAN gateway, all networks locally-attached to a firewall interface, the configured DNS servers, VPN addresses and Virtual IP addresses. Additional HOME_NET networks may be created on the IP LISTS tab, and then return to this tab to assign them to the Snort interface by selecting the appropriate list in the drop-down selector. View the contents of the selected list by clicking the **View List** button.

External Net: selects the network will use as the EXTERNAL_NET variable. Default is the recommended choice and contains all networks not included in HOME_NET. Create additional EXTERNAL_NET networks on the IP LISTS tab, and then return to this tab to assign them to the Snort interface by selecting the appropriate list in the drop-down selector.

Pass List: selects the networks and IP addresses that Snort will never block. These represent "trusted hosts". Even if a trusted host generates a Snort alert, it will not be blocked if the IP address is on a Pass List. The default Pass List contains the same addresses as HOME_NET. Create additional pass lists on the IP LISTS tab, and then return to this tab to assign them to the Snort interface by selecting the appropriate list in the drop-down selector. Snort must be restarted on the interface when making changes to the Pass List. View the contents of the selected list by clicking the **View List** button.

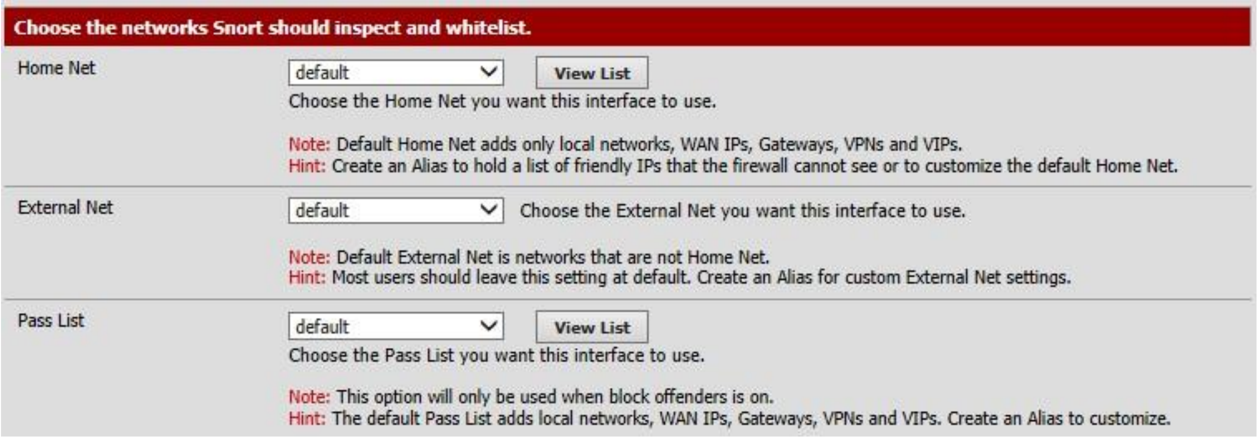

Choose the networks Snort should inspect and whitelist.

Home Net: default

Choose the Home Net you want this interface to use.

Note: Default Home Net adds only local networks, WAN IPs, Gateways, VPNs and VIPs.
Hint: Create an Alias to hold a list of friendly IPs that the firewall cannot see or to customize the default Home Net.

External Net: default
Choose the External Net you want this interface to use.

Note: Default External Net is networks that are not Home Net.
Hint: Most users should leave this setting at default. Create an Alias for custom External Net settings.

Pass List: default

Choose the Pass List you want this interface to use.

Note: This option will only be used when block offenders is on.
Hint: The default Pass List adds local networks, WAN IPs, Gateways, VPNs and VIPs. Create an Alias to customize.

Choose a suppression or filtering file if desired

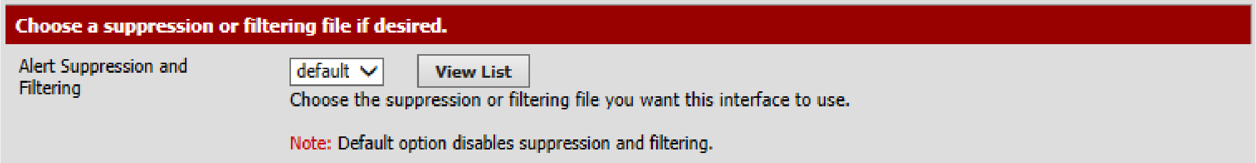

Choose a suppression or filtering file if desired.

Alert Suppression and Filtering: default

Choose the suppression or filtering file you want this interface to use.

Note: Default option disables suppression and filtering.

Snort interface Global Settings

This tab is used to enable rule set packages for download, configure the rules package update interval and start time, configure Snort logging directory size limits and determine whether Snort settings are saved when the package is removed from the system.

Rules Update Settings

Use the **Update Interval:** drop-down selector to choose the periodicity for checking for updates to the enabled rules packages. When any value other than NEVER is selected, the **Update Start Time** textbox is available for entering a start time in 24-hour format using hours and minutes only.

In most cases every 12 hours is a good choice. The update start time can be customized if desired. Enter the time as hours and minutes in 24-hour time format. The default start time is 3 minutes past midnight local time. So with a 12-hour update interval selected, Snort will check web sites at 3 minutes past midnight and 3 minutes past noon each day for any posted rule package updates.

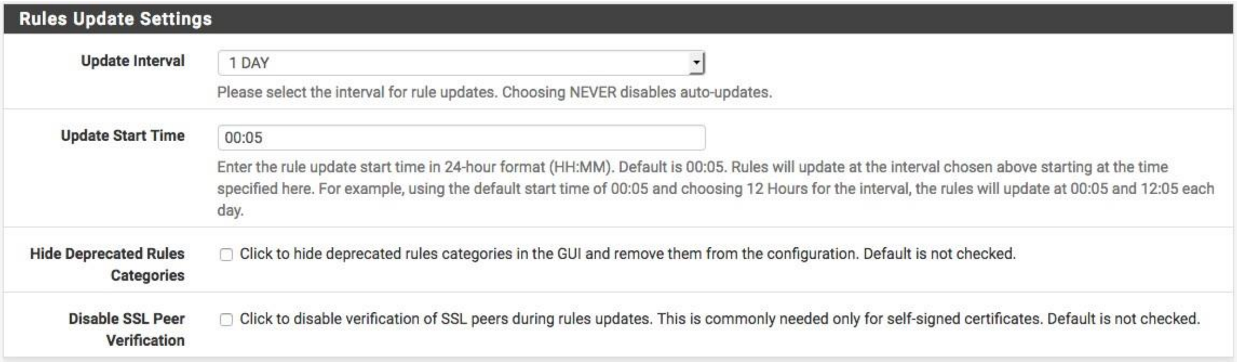

Rules Update Settings

Update Interval	1 DAY Please select the interval for rule updates. Choosing NEVER disables auto-updates.
Update Start Time	00:05 Enter the rule update start time in 24-hour format (HH:MM). Default is 00:05. Rules will update at the interval chosen above starting at the time specified here. For example, using the default start time of 00:05 and choosing 12 Hours for the interval, the rules will update at 00:05 and 12:05 each day.
Hide Deprecated Rules Categories	□ Click to hide deprecated rules categories in the GUI and remove them from the configuration. Default is not checked.
Disable SSL Peer Verification	□ Click to disable verification of SSL peers during rules updates. This is commonly needed only for self-signed certificates. Default is not checked.

General Settings

The **Log Directory Size Limit**, when enabled, sets an absolute hard upper limit on the total size of the Snort logging sub-directory in ***/var/log/snort***. This can prevent Snort from filling up the **/var** volume on the firewall. When the Snort logging directory size (the total size of all files within the Snort log directory tree) exceed the value set, all files are automatically pruned (deleted) and the Snort process is signaled to soft-restart and resynchronize logging. The default size limit is 20% of the available space on the volume. This may be overridden by setting a value in megabytes (MB) in the textbox provided.

Remove Blocked Hosts Interval: controls how long Snort-blocked IP addresses must be inactive before being cleared. Once per interval specified, Snort executes a cron job that tests all the IP addresses it has inserted into the firewall's block table for activity. IP addresses that have had no further network activity within the time specified are removed from the block table.

Remove Blocked Hosts After Deinstall: determines whether or not Snort-blocked IP addresses are automatically removed when the Snort package is uninstalled.

Remove Snort Log Files After Deinstall: determines whether or not log files generated by Snort are retained or removed when the Snort package is removed.

Keep Snort Settings After Deinstall: controls whether the Snort configuration is retained when the Snort package is removed.

Snort Interfaces

The **Snort Interfaces** tab is where one can add, edit or delete a Snort instance from a physical network interface. A snort instance can also manually started and stopped. If *Barnyard2* is configured on an interface, it can also be started or stopped.

The green icon indicates a running Snort process for the interface. To stop a running process, click the ✓ icon and wait for it to change to ✖

To add a new Snort configuration for an interface,

click **Add**. To edit an existing Snort configuration,

click edit icon.

To delete an existing Snort configuration, click the checkbox on the left of the interface to select it, then click **Delete**. A prompt for confirmation will appear before deleting the chosen interface. Multiple interfaces may be selected and deleted at once.

Managing Snort IP Address Lists

Use this tab to manage the IP lists files for the IP Reputation preprocessor. IP lists are text-format files containing one IP address or network (expressed in CIDR notation) per line.

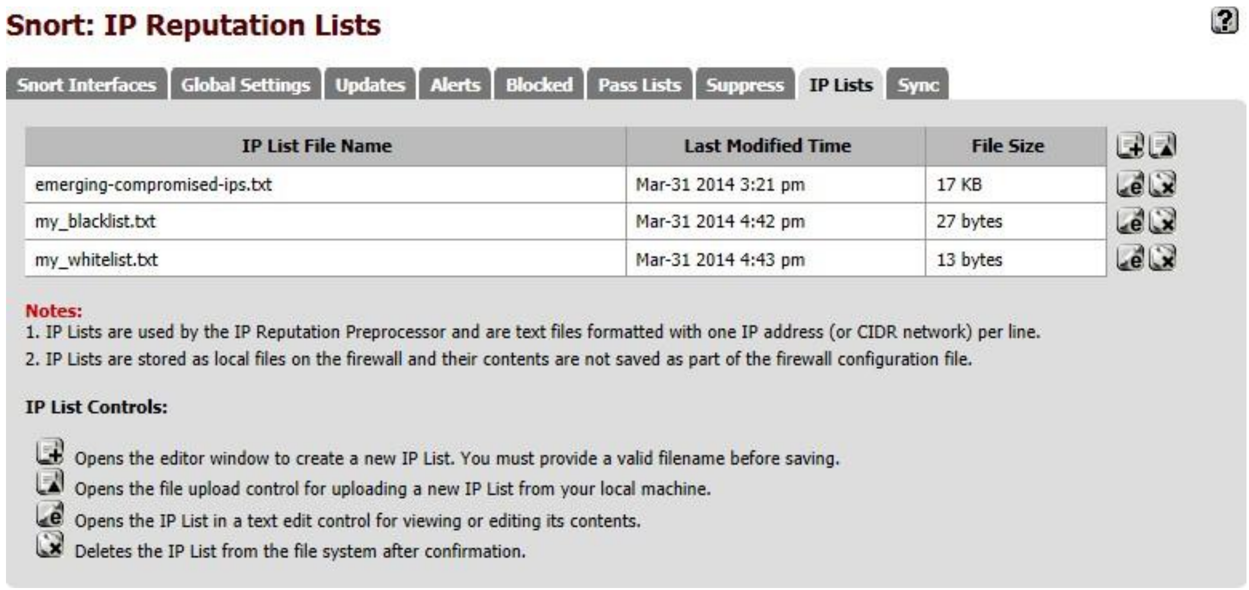

Snort: IP Reputation Lists

IP List File Name	Last Modified Time	File Size
emerging-compromised-ips.txt	Mar-31 2014 3:21 pm	17 KB
my_blacklist.txt	Mar-31 2014 4:42 pm	27 bytes
my_whitelist.txt	Mar-31 2014 4:43 pm	13 bytes

Notes:
1. IP Lists are used by the IP Reputation Preprocessor and are text files formatted with one IP address (or CIDR network) per line.
2. IP Lists are stored as local files on the firewall and their contents are not saved as part of the firewall configuration file.

IP List Controls:

Opens the editor window to create a new IP List. You must provide a valid filename before saving.
Opens the file upload control for uploading a new IP List from your local machine.
Opens the IP List in a text edit control for viewing or editing its contents.
Deletes the IP List from the file system after confirmation.

To upload an IP list file to the firewall, click the ⬆ icon to open the file upload dialog as shown below. Browse to the file on the local machine using the **BROWSE** button, then click the

UPLOAD button to upload the file to the firewall for use by the IP Reputation preprocessor in Snort.

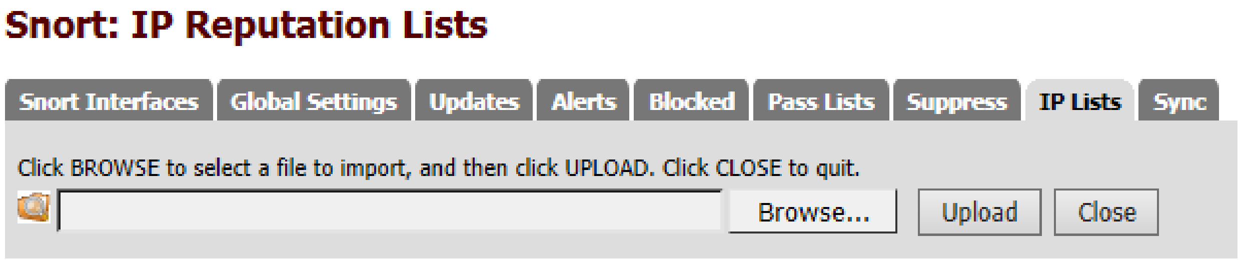

Click BROWSE to select a file to import, and then click UPLOAD. Click CLOSE to quit.

To create a new IP list, click the ➕ icon. To edit an existing IP list, click the ✏ icon beside the list to edit. Click SAVE when finished to save changes to the list, or CANCEL to abandon any changes.

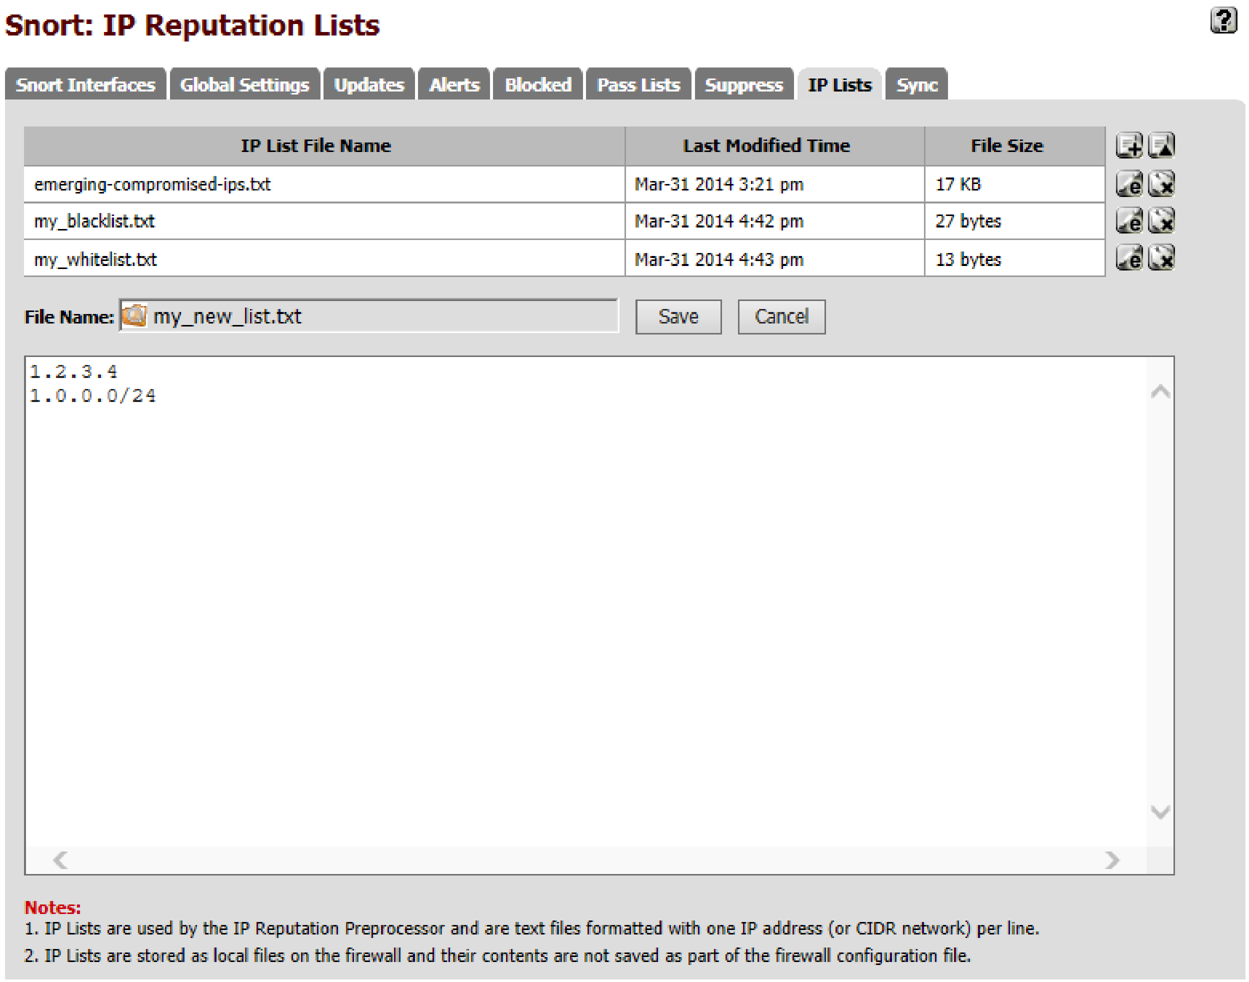

IP List File Name	Last Modified Time	File Size
emerging-compromised-ips.txt	Mar-31 2014 3:21 pm	17 KB
my_blacklist.txt	Mar-31 2014 4:42 pm	27 bytes
my_whitelist.txt	Mar-31 2014 4:43 pm	13 bytes

File Name: my_new_list.txt

```
1.2.3.4
1.0.0.0/24
```

Notes:
1. IP Lists are used by the IP Reputation Preprocessor and are text files formatted with one IP address (or CIDR network) per line.
2. IP Lists are stored as local files on the firewall and their contents are not saved as part of the firewall configuration file.

Snort IP Address Reputation Preprocessor

This tab allows configuration of the parameters specific to the IP Reputation preprocessor on the interface. It also allows the assignment of blacklist and whitelist files of IP addresses to the interface.

The available fields are:

Enable: when checked, the IP Reputation preprocessor is active on this Snort instance.

Memory Cap: sets the amount of system memory in megabytes (MB) to reserve for storage of the IP lists associated with this preprocessor. The default is 500 MB and should be sufficient for most installations.

Scan Local: when checked, Snort will include RFC 1918 IP addresses ranges when comparing IP addresses to the blacklists and whitelists. If an RFC 1918 IP addresses is in the whitelist files, or some are blacklist files, then this option should be enabled. The default is disabled.

Nested IP: this tells Snort which IP address to compare to the IP lists in the whitelist and blacklist files when there is IP encapsulation. The default is **Inner**.

Priority: instructs Snort which IP list has priority when the source and destination IP addresses of a packet are each on separate IP lists. For example, if the source IP address is on a blacklist while the destination IP address is on a whitelist, this option tells Snort whether to block the traffic if blacklist has priority, or pass the traffic if whitelist has priority.

Whitelist Meaning: this tells Snort what action to take with whitelisted IP addresses. The two options are **Un-black** and **Trust**. When set to **Un-black**, a blacklisted IP which is listed in the whitelist is not immediately blocked. Instead it is routed through the Snort detection engine for normal inspection. If it generates no alerts, the traffic is allowed. If the inspection results in a Snort alert for the traffic, it will be blocked.

When set to **Trust**, any IP address on the whitelist (including any that may also be on a blacklist) is immediately allowed to pass with no further inspection. Caution should be exercised when using the Trust mode of operation to insure the IP addresses on the whitelist are in fact trustworthy.

The ➕ and 🗑 icons at the bottom of the page are used to assign or remove blacklist and whitelist files to or from the interface.

Click the ➕ icon to open a file selection dialog. Choose an IP list file from the list by clicking on the name.

Snort Pass Lists

Pass Lists are lists of IP addresses that Snort should never block. Pass lists can be created and managed on the **Pass Lists** tab. When an IP address is listed on a Pass List, Snort will never insert a block on that address even when malicious traffic is detected.

To create a new Pass List, click the ➕ icon. To edit an existing Pass List, click the ✏ icon.

To delete a Pass List, click the 🗑 icon. Note that a Pass List cannot be deleted if it is currently assigned to one or more Snort interfaces.

A default Pass List is automatically generated by Snort for every interface, and this default list is used when no other list is specified. Assign Pass Lists to an interface on the **Interface Settings** tab.

Customized Pass Lists can be created and then assigned to an interface. This might be needed when trusted external hosts are needed that are not located on networks directly connected to the firewall. To add external hosts in this manner, first create an Alias under **Firewall** > **Aliases** and then assign that alias to the **Assigned Aliases:** field. In the

example shown below, the alias "*Friendly_ext_hosts*" has been assigned. This alias would contain the IP addresses of the trusted external hosts.

When creating a custom Pass List, leave all the auto-generated IP addresses checked in the **Add auto-generated IP addresses** section. Not selecting the checkboxes in this section can lead to blocking of critical addresses including the firewall interfaces themselves. This could result in being locked out of the firewall over the network! Only uncheck boxes in this section when a valid need is present.

Click the **ALIASES** button to open a window showing previously defined aliases for selection. Remember to click
SAVE to save changes.

Snort Suppression Lists

Alert Thresholding and Suppression

Suppression Lists allow control over the alerts generated by Snort rules. When an alert is suppressed, then Snort no longer logs an alert entry (or blocks the IP address if block offenders is enabled) when a particular rule fires. Snort still inspects all network traffic against the rule, but even when traffic matches the rule signature, no alert will be generated. This is different from disabling a rule. When a rule is disabled, Snort no longer tries to match it to any network traffic. Suppressing a rule might be done in lieu of disabling the rule to stop alerts based on either the source or destination IP. For example, to suppress the alert when traffic from a particular trusted IP address is the source. If any other IP is the source or destination of the traffic, the rule may still be desired. To eliminate all alerts from the rule, then it is more efficient to simply disable the rule rather than to suppress it. Disabling the rule will remove it from the list of match rules in Snort and therefore makes for less work Snort has to do.

On the Suppress List Edit page, suppress lists may be manually added or edited. It is usually easier and faster to add suppress list entries by clicking the ➕ icons shown with the alert entries on the **Alerts** tab. Remember to click the **SAVE** button to save changes when manually editing Suppress List entries.

Lists with comments are easier to manipulate and fine tune. Neither screen shot shows IP address suffix in a suppress entry.

24.7Status Traffic Totals

This package displays different ways to view the traffic usage generated by the network traffic monitoring tool vnStat.

CHAPTER

TWENTYFIVE

CELLULAR WIRELESS

AZTCO-FW® software can use a supported cellular modem (3G/4G/LTE) as a WAN interface for connectivity. This can be used as a sole means of connectivity or as a backup.

25.1Configuring 3G modems

To configure a 3G modem in AZTCO-FW® software on a current supported release, plug in a *Known Working Modem* and log into the web interface to begin configuration.

25.1.1Configuring PPP

Browse to **Interfaces > Assignments**, and click the **PPPs** tab. Click **+** on that screen. In

the **Link Type** drop down, select *PPP*.

In the **Link interface(s)** box, the list will be populated with serial ports on the system. Select the port for the modem. A modem may list several serial ports. Typically, it is the last one, but may require some trial and error. Future versions of AZTCO-FW software may properly auto-detect modems but that has historically been a source of problems.

Optionally fill in a **Description**, which will be used to reference this PPP configuration in other parts of the web interface.

Under **Service Provider**, select the **Country**. The **Provider** list for that country will appear, then select the provider of the card. Then in the **Plan** drop down, select the plan. This should adequately fill in all the PPP details needed for the connection. Click **Save**.

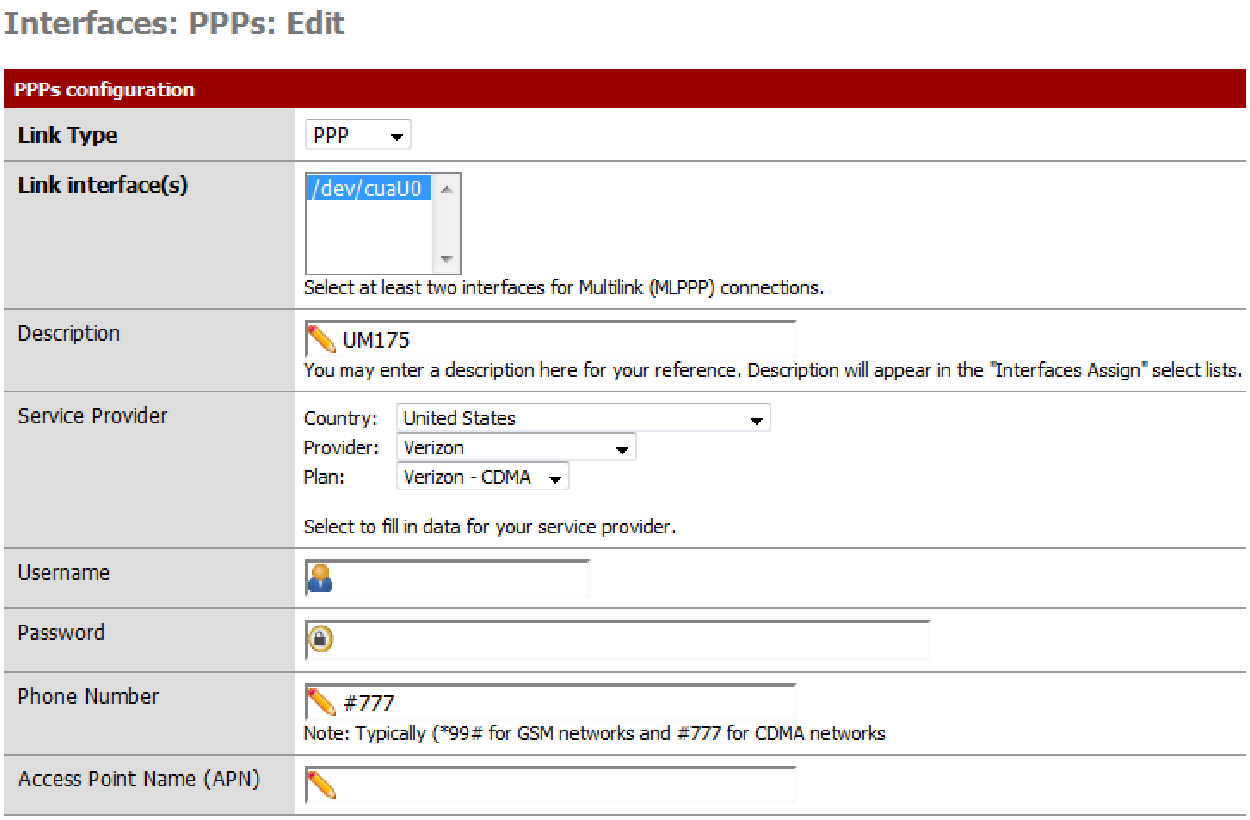

Interfaces: PPPs: Edit

PPPs configuration	
Link Type	PPP
Link interface(s)	/dev/cuaU0 Select at least two interfaces for Multilink (MLPPP) connections.
Description	UM175 You may enter a description here for your reference. Description will appear in the "Interfaces Assign" select lists.
Service Provider	Country: United States Provider: Verizon Plan: Verizon - CDMA Select to fill in data for your service provider.
Username	
Password	
Phone Number	#777 Note: Typically (*99# for GSM networks and #777 for CDMA networks
Access Point Name (APN)	

At the PPPs screen, the newly created PPP interface will be listed.

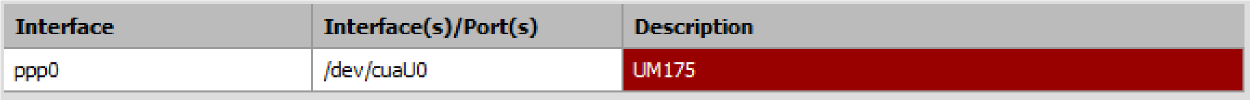

Interface	Interface(s)/Port(s)	Description
ppp0	/dev/cuaU0	UM175

25.1.2Assigning the PPP Interface

Next the PPP interface must be created. Click the **Interface assignments** tab, and click the **+** to add a new interface. Select the PPP interface, click **Save**, then **Apply changes**.

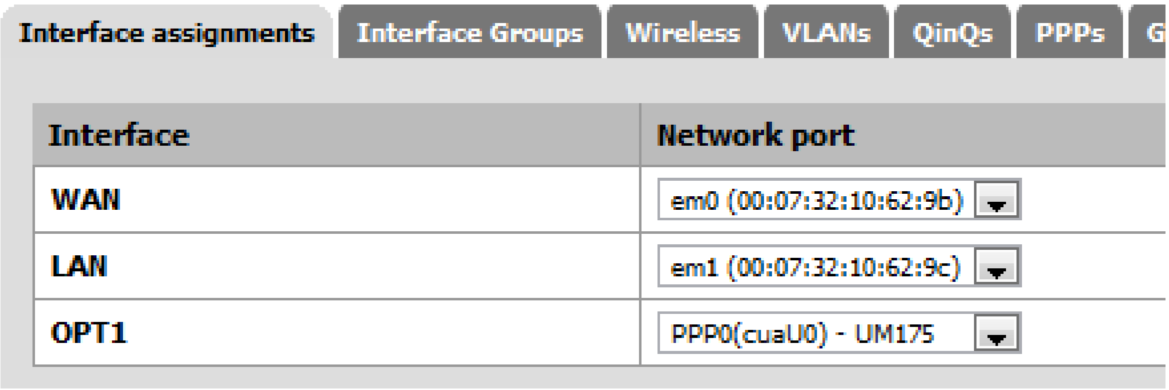

Interface	Network port
WAN	em0 (00:07:32:10:62:9b)
LAN	em1 (00:07:32:10:62:9c)
OPT1	PPP0(cuaU0) - UM175

25.1.3Enable the PPP Interface

Now browse to **Interfaces > OPT1** (or the interface name shown for the PPP interface when it was assigned above). Check the **Enable Interface** box, rename it if desired, and click **Save**. Do not change anything else on the page. Then click **Apply changes**.

25.1.4Check the interface status

Browse to **Status > Interfaces** to check the status of the newly created and assigned PPP interface. If it does not show connected, check the logs under **Status > System logs**, **PPP** tab to see why its connection isfailing.

Note that some connection problems are a lack of signal. If the 3G modem is in a location with poor reception, such as an equipment room, datacenter, rack, etc, then it may be advisable to use a longer USB cable and/or Antenna to achieve a better signal.

END OF DOCUMENTATION