Edge Threat Management (ETM) Micro Edge User Manual

Headquarters	Support	Sales
5453 Great America Parkway Santa Clara, CA 95054 USA +1-408-547-5500	+1-408-547-5502 +1-866-476-0000	+1-408-547-5501 +1-866-497-0000
www.arista.com/en/	support@arista.com	sales@arista.com

Contents

Getting Started

Contents

1.1   Hardware Requirements

This section provides information on hardware requirements for running NG Firewall.

Hardware Requirements

- NG Firewall must be installed onto a dedicated machine with at least two Network Interface Controller (NICs).
- NG Firewall installs to the hard drive of a PC, **erasing all data on that drive in the process**. Be aware of this before starting the installation.

Purchasing an NG Firewall Appliance

You can purchase a server directly from us, pre-installed with our software. Please see our appliances page for more information.

Using your own Hardware

One of the great things about NG Firewall is that we're software based, meaning you can install it to any desktop or server PC that fits the bill. When assembling an NG Firewall server, you must account for several variables - the hardware you're installing onto, the number of users as well as the workload of those users. The table below provides some recommended hardware specs by network size.

Remember that these are only guidelines: each network is different, each user is different, each configuration is different.

Table 1: Hardware Recommendations Table

Resource	Processor	Memory	Hard Drive	Notes
Minimum Spec	1 core	2 GB	40 GB	Platform only (no apps)
1-50 devices	2 cores	4 GB	40 GB	Arista ETM z4 equivalent
51-150 devices	4 cores	8 GB	80 GB	
151-500 devices	4 or more cores	16 GB	160 GB	

Resource	Processor	Memory	Hard Drive	Notes
501-1500 devices	4 or more cores	16 GB	250 GB	Arista ETM z12 & Q12 equivalent
1501-5000 devices	6 or more cores	32 GB	500 GB	Arista ETM z20 & Q20 equivalent

Hardware Compatibility

If you are using NG Firewall on your own hardware you will need to do the research to determine if the hardware is compatible. Generally speaking, if a particular piece of hardware is known to work well in recent versions of Linux, then it has a good chance of working with NG Firewall.

To know for sure, you will need to order the equipment and install NGFW on it. If it works - it is compatible. If it does not work, you can try tweaking BIOS settings and possibly swapping parts and/or monitor and disk configurations. If it still does not work, the hardware is likely incompatible.

For users that do not want to go through this process, buying an appliance is suggested as that is the only hardware we can guarantee will work and the only hardware we support.

We do not suggest USB NICs. In our experience, they perform poorly and are not reliable.

1.2   NG Firewall Installation

NG Firewall is NGFW/UTM software, bringing together everything your network needs to stay healthy in one box: web content and spam filtering, virus scanning, VPN connectivity, multi-WAN failover capability, and much more.

This guide will be a quick primer on getting your NG Firewall installed, up and running. Hopefully, it will also answer some common configuration questions without causing too much confusion. If you already have NG Firewall in your network, you can skip to any relevant section and read from there. If you're new to NG Firewall, we recommend reading this section to help familiarize yourself with the software and how it works - it will probably save you a headache or two later on.

What is NG Firewall?

We strive to make deployment and administration easy with a friendly web-based GUI to help you monitor and filter traffic on your network. NG Firewall provides a suite of applications free of charge with the option of subscribing to additional applications as best suits your organization - our website has a full list of features. Current pricing for paid applications, packages, and appliances can be found in the store.

Deploying NG Firewall

NG Firewall is available in the following deployment options:

- **Cloud Appliance**: A virtual appliance available for Amazon Web Services or Microsoft Azure. Learn more about the AWS and Azure public cloud appliances here.
- **Virtual Appliance**: A virtual appliance optimized for VMware deployments in private cloud infrastructure. You can download the virtual and software appliances from ETM Dashboard. The virtual appliance is available as an OVA formatted file. See NG Firewall Virtual Appliance on VMware for installation details.
- **Hardware Appliance**: An Arista Edge Threat Management network appliance with NG Firewall preinstalled. Learn more about the zSeries appliances here.
- **Software Appliance**: An installable version of NG Firewall for most x86 based devices. The software appliance is available as an ISO formatted file that you can image to a USB drive. See Creating a bootable USB installer for imaging instructions.

Installing the NG Firewall Software Appliance

The software appliance method installs to the primary storage of a device, **erasing all data on that drive in the process**. Be aware of this before starting the installation. Also, note that NG Firewall **requires** at least two NICs to be installed **before** you start the installation.

Most users install NG Firewall on the server before the server is placed in-line on their network. To do this, plug one interface of your NG Firewall into your network as you would any other computer, then start the installer. This ensures that NG Firewall will have access to the internet during installation.

Power down the server, insert the ISO or USB installer, and power on the server. Ensure the boot options are set to boot from the inserted CD or USB media. Once the installation has started, follow the directions on the screen to complete the installation process.

During the installation, you may need to answer a few questions, for example, to confirm writing to the storage device. If you encounter issues while installing NG Firewall onto your server, read the Troubleshooting Server Installation.

UEFI Installation

As of release *16.0*, NG Firewall can be installed via BIOS or UEFI. When booting via CD or USB, the installer automatically detects whether it was booted via BIOS or UEFI and tweaks the install process accordingly. To tell whether the installer was booted via BIOS or UEFI, check the installer's menu title. When booted via BIOS, the installer menu title will be "NG Firewall installer boot menu". When booted via UEFI, the installer menu title will be "NG Firewall UEFI Installer".

Serial Console Installation

As of version *16.5*, you can install and manage NG Firewall via a serial console port. This is useful if your device does not have video output and supports serial management. This method uses a dedicated ISO installer that you must download. Your system must be configured to use **S0** as the serial interface and a baud rate of **1115200**.

Account Registration

NG Firewall will prompt you to sign in or register a new account in ETM Dashboard. Registration is required to install any applications and takes only a second.

Registration has the following benefits:

- Install free or paid applications on your NG Firewall.
- Manage your licenses, renewals, servers, and contact info all from one dashboard.
- Easily transfer licenses between servers.

If you signed in with an existing account, the system will check for any unused subscriptions in your account and ask if you would like to apply them to this system.

Once you have completed the process, continue with the steps below. Your account can always be accessed by visiting https://edge.arista.com or clicking *My Account* in the lower left-hand corner of the UI.

Install Applications

Installing applications is covered in the NG Firewall User Guide. It is recommended to finish reading this section and get everything working before configuring/tuning the application settings.

Configure Other Subnets

NG Firewall will route all traffic according to its routing table, even when installed as a *Transparent Bridge.* This means NG Firewall must have the proper routing table for all subnets on your network.

If you have other subnets on the network aside from those configured in the Setup Wizard, you will need to configure NG Firewall to know about these networks. For example, if you are running as a bridge with NG Firewall having an address **192.168.1.2** with a netmask **255.255.255.0**, but you also have a **192.168.20.*** network and also a **10.0.*.*** network, you will need to tell NG Firewall where to reach these hosts.

There are several ways to do this:

- Add a route in **Config** > **Network** > **Routes** telling NG Firewall how to reach those subnets. If **10.0.*.*** is local on Internal, then you simply need to create a **10.0.0.0/16** route to "Local on Internal." If **10.0.*.*** lives behind another router on your network, like **192.168.1.100** then you must add a route to send all **10.0.0.0/16** traffic to **192.168.1.100**.
- Add an alias on the appropriate interface. In **Config** > **Network** > **Interface**, click **Edit** on the appropriate interface and add an alias IP. This tells NG Firewall that this IP range is local and can be reached locally on that interface. It also provides NG Firewall a local address on those subnets should any of those clients need to reach NG Firewall using a local IP.

Each subnet on your network must be configured so NG Firewall knows how to reach them. The "Ping Test" in **Config** > **Network** > **Troubleshooting** can be used to verify that NG Firewall can reach the configured subnets.

More in depth information about how the network is configured is found in Network Configuration.

Configure Other Interfaces

In the setup wizard, you configured both the Internal and External interfaces. If you have more than 2 interfaces, the 3rd and beyond are *Disabled* by default.

If you plan to use them, they must be configured, and it is suggested to choose a name reflecting its use.

Common uses include:

Additional WAN interfaces (if you have multiple internet connections) for failover/balancing

To do this just configure it as a WAN interface with the ISP's provided values. Read more about WAN Failover and WAN Balancer for more information about failover/balancing.

Other internal networks

To do this, just configure it as a non-WAN interface with a static internal IP. For example, if you used **192.168.1.1/24** on your internal, you could use **192.168.2.1/24** on your 3rd interface. This is useful on larger networks, for guest networks, for wireless networks.

Public segment for public servers (DMZ)

If you have servers with public addresses, you can stick them on the additional interface(s) and bridge those interfaces to your WAN. Then, configure them with IPs on the same subnet as the WAN interface.

Additional NICs for existing networks

If you want additional NICs for you Internal (for example), you can bridge the 3rd interface to your Internal and plug in additional internal machines to that NIC. This behaves similarly to a switch, but traffic going through the NG Firewall to reach other internal hosts is scanned by the apps.

Configuring a WiFi interface

If your hardware platform includes a supported WiFi adapter, you can configure your WiFi interface. Be sure to select the appropriate Regulatory Country option for your country.

More in depth information about how the network is configured is found in Network Configuration.

Email

Some NG Firewall applications and functions rely on sending emails like reports and spam quarantine digests. Email sending is configured in **Config** > **Email**. By default, email will be sent directly using DNS MX records like a mail server. However, some ISPs and networks block port *25* to prevent spam, and in this case, you must configure a SMTP relay (and the appropriate authorization credentials if required).

Hostname

You can configure the hostname (and domain) for the NG Firewall server in **Config** > **Network** > **Hostname**.

Port Forward Rules

If NG Firewall is installed as a router and has internal servers with services that need to be publicly accessible, you need to configure port forward rules to forward that traffic to the appropriate server. You can configure port forward rules in **Config** > **Network** > **Port Forward Rules**.

Bypass Rules

Unlike many next-generation firewalls, NG Firewall scans *All* TCP and UDP traffic on all ports at the application layer by default, except for VoIP traffic. This is ideal for most deployments, but if you are running a very large (1000s of users) network, it probably makes sense to bypass traffic that you are not interested in scanning. Traffic can be bypassed in **Config** > **Network** > **Bypass Rules**. More is described in the Network documentation.

Public Address

If you use OpenVPN, quarantine or, other publicly accessible services on NG Firewall, you may wish to configure the "public address" of NG Firewall so that it sends the appropriate URL to remote users. Public Address is configured in **Config** > **Administration**.

External Administration

If you'd like to be able to administer NG Firewall via HTTPS remotely, you will need to enable HTTPS access on WAN interfaces in the Filter Rules .

Installing NG Firewall on the Network

At this point, NG Firewall should be ready to drop into the network if it is not already in place.

If NG Firewall is configured in bridge mode, an easy way to test is to install it with only one or a few computers behind it - plug the External interface into your network, then plug a switch with a few computers into the Internal interface so they must go through NG Firewall. Only those computers will be filtered, allowing you to test without disturbing the rest of your network.

If you are running as a *Transparent Bridge,* verify that NG Firewall is not plugged in backwards by unplugging the network cables one at a time and looking at the green lights in **Config** > **Network** > **Interfaces**. If NG Firewall is configured as a bridge and plugged in backwards, it will pass traffic, but some functionality will not work correctly. NG Firewall also provides Administration Notifications, which will bring this to your attention so you can fix it.

- NG Firewall is designed to drop into your network with minimum disruption. When testing we recommend putting the system in place, keeping most defaults unless you're having problems. This way, you can get a feel for how it works before making major changes that may affect system operation.

Using NG Firewall

The next step is installing the applications and configuring NG Firewall to meet your needs. The NG Firewall User Guide provides in depth documentation of the various functions of NG Firewall and the applications.

Welcome to NG Firewall!

Related information
https://edge.arista.com

1.2.1   Setup Wizard

The *Setup Wizard* will open automatically when NG Firewall first boots.

If you do not have a keyboard/mouse/video connected to the NG Firewall server, the Setup Wizard can be reached by plugging into a DHCP-configured laptop into the internal interface and opening a browser to **http://192.168.2.1/**.

Once installed, the setup wizard can be repeated at any time and found in the NG Firewall GUI at **Config > System > Setup Wizard**.

Welcome Page

For versions *16.3* and newer, the Setup Wizard begins with a welcome page. Choose to create an ETM Dashboard account or log in with an existing one to get started. Your ETM Dashboard account is free and is necessary to activate a trial or complete license on the device. Your account is also linked to ETM Dashboard, enabling you to remotely manage your Arista Edge Threat Management appliances.

By logging in or creating your ETM Dashboard account, the Add Appliance wizard opens automatically and includes the UID of your appliance. The Add Appliance wizard guides you through the remainder of the setup steps for your new NG Firewall appliance. See Adding Appliances to ETM Dashboard for more details.

To continue, you must log in using your Command Center account. If you do not have one, you can create a free account.

LOG IN

CREATE ACCOUNT
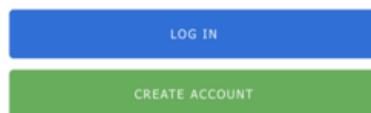
If your NG Firewall device is not connected to the Internet or requires specific configuration to connect, the wizard allows you to Configure the Internet Connection. If you are unable to connect to the Internet, you can continue with the local setup wizard by following these instructions: Offline Setup Wizard.

The next steps include installing the desired apps and possibly tuning the configuration of your NG Firewall.

1.2.2   Offline Setup Wizard

If your NG Firewall appliance cannot connect to the Internet and your appliance is not configured, the local Administration Interface presents you with a Setup Wizard to configure essential parameters. If your appliance is online and connected to the Internet, the configuration of essential parameters is managed through ETM Dashboard.

The steps below explain how to navigate through the local Setup Wizard if your appliance is offline.

Language Selection

Before you begin the setup wizard, select your preferred language.

Figure 1: ETM Language Setup

Setup Wizard - Welcome

The next screen simply welcomes you to the Setup Wizard. Click next to continue.
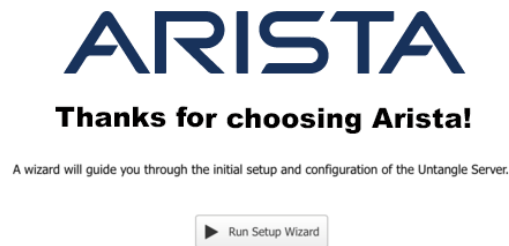
Figure 2: Setup Wizard Step0

Setup Wizard - Step 1 - Configure the Server

The first step is to set a password for the administrator account and select a timezone. You can also set the admin email to receive alerts and reports. Select the installation type, which closest matches your deployment.
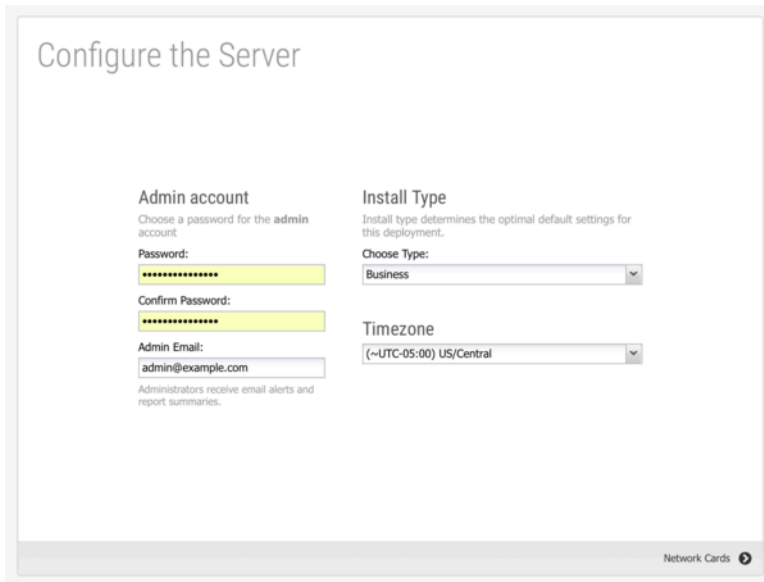
Figure 3: Setup Wizard Step1

Setup Wizard - Step 2 - Identify Network Cards

The second step shows you the network cards. If this is an appliance from Edge Threat Management, you can simply continue to the next step. If this is a custom server, verify that the physical network cards are mapped to the correct (desired) interface. You can verify connectivity by disconnecting or connecting the cable to the physical interface. The status icon changes immediately between grey or green to show the link state.
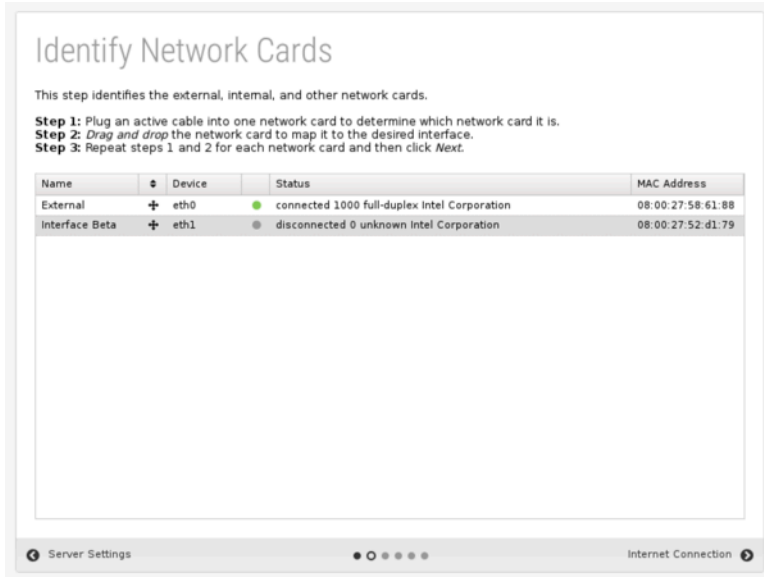
Figure 4: Setup Wizard Step2

Setup Wizard - Step 3 - Configure The Internet Connection

The third step configures your External (WAN) interface.
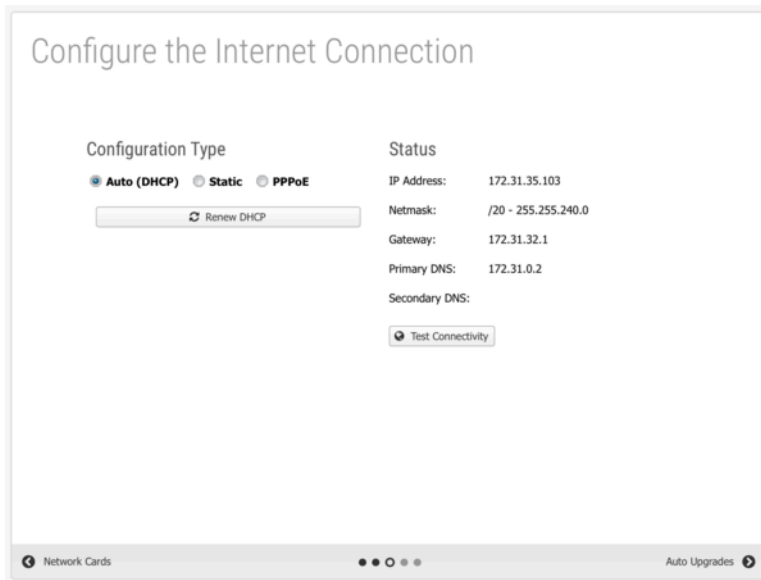
Figure 5: Setup Wizard Step3

The default selection is *Auto (DHCP)*. The automatically assigned address is displayed if an address was successfully acquired. Otherwise, click **Renew DHCP** to acquire an IP address. Click **Test Connectivity** to verify Internet access.

If your Internet connection requires a static IP address or uses PPPoE, select the appropriate option and enter the parameters assigned by your Internet Service Provider.

Setup Wizard - Step 4 - Internal Network Interface

The fourth step will configure your "Internal" interface (and DHCP server and NAT configuration.) There are two choices.
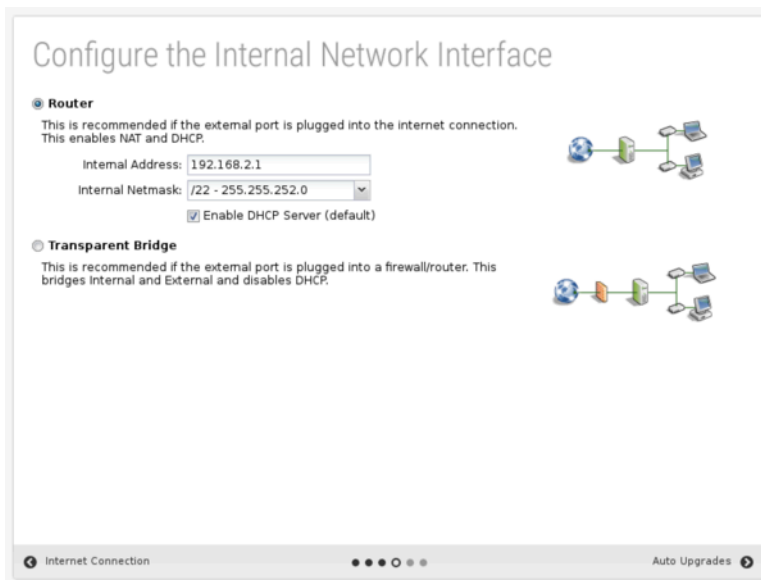
Figure 6: Setup Wizard Step4

You can configure the internal interface with private static IP address such as *192.168.2.1* and enable DHCP serving and Network Address Translation (NAT) so all internal machines will have private addresses and share one public IP. This is commonly referred to as *Router* mode.

You can also configure the internal interface to be bridged to the external. In this mode, the internal interface does not have its address and simply shares the External address. This is commonly referred to as *Transparent Bridge* mode.

Router

In **Router** mode, NG Firewall will be the edge device on your network and serve as a router and firewall. In this case you'll need to set up your External and Internal interfaces correctly for traffic to flow, which should have been done while installing.

Transparent Bridge

In **Transparent Bridge** mode, NG Firewall is installed behind an existing firewall and sits between your existing firewall and main switch. When in Bridge mode NG Firewall is transparent, meaning you won't need to change the default gateway of the computers on your network or the routes on your firewall - just put the NG Firewall between your firewall and main switch, and that's it. You do not change the configuration of existing clients or the existing firewall!

Setup Wizard - Step 5 - Configure Automatic Upgrade Settings

In the fifth step, *Automatic Upgrades* are configured. If Automatic Upgrades is enabled, NG Firewall automatically checks for new versions and performs the upgrade between 1 am and 2 am every morning. You can adjust the upgrade schedule after the setup is complete from the Upgrade.

This step also includes an option to manage the appliance from ETM Dashboard.
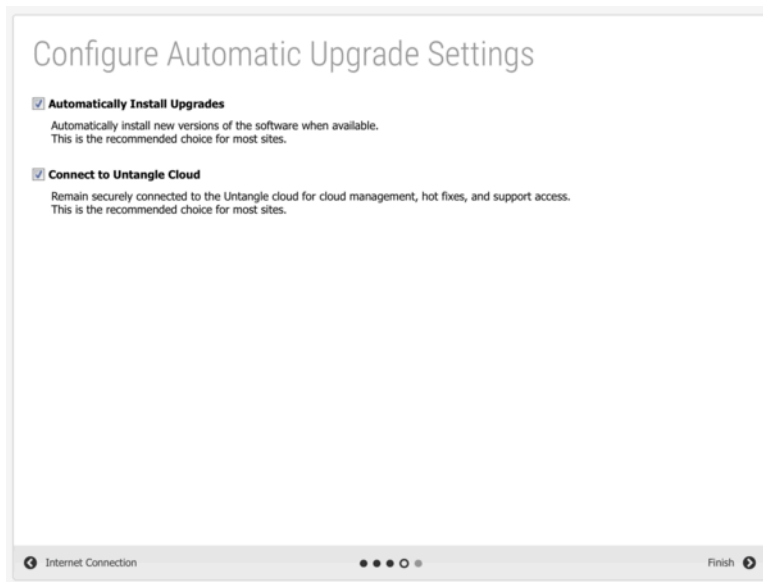
Figure 7: Setup Wizard Step5

Setup Wizard - Finished

That's it!

Figure 8: Setup Wizard Step6

1.2.3 Upgrade

Upgrade allows the server to upgrade and contains upgrade-related settings.

Upgrade Settings

Upgrades show the currently available upgrades, if any. If upgrades are available, an upgrade can be started by pressing the *Upgrade* button at the top under **Status**.

To see changes, see the Changelogs.
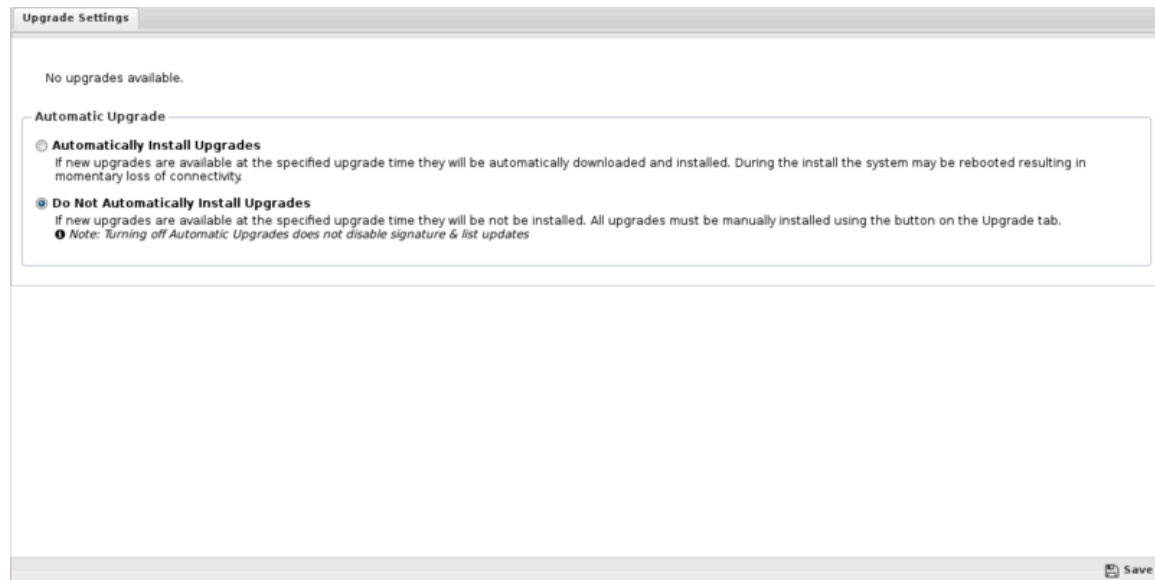
Figure 9: Upgrade Settings

After the upgrade begins, it will download the new packages (which may take some time), and then the upgrades will be applied. Do not reboot or power off the server during the upgrade.

If **Automatically Install Upgrades** is checked, NG Firewall will automatically check for new versions and upgrade if available.

Automatic Upgrade Schedule configures when NG Firewall will automatically upgrade if upgrades are available. NG Firewall will automatically upgrade at the specified time on the days of the week that are checked.

1.2.4 Troubleshooting Server Installation

Video Issues

Occasionally Arista can not correctly detect video card/monitor settings to display successfully on the monitor. This can happen in several ways:

- The monitor flashes and then displays a black screen with a message or login prompt
- The monitor just displays noise after the bootup is complete
- The monitor displays correctly, but the screen is much too big, requiring scrolling with the mouse.

Figure 10: Monitor Problem

Things to try:

- Restart the server and select a different boot mode in the bootup kernel selection menu.
- Try various BIOS settings that may affect the video.
- Try another monitor and reboot after switching. The monitor should be plugged in before powering on Arista.
- If you are using a KVM (keyboard-video-monitor switch), remove it and connect the peripherals directly.
- Try another video card.
- Re-burn the CD/ISO slower, or re-create the USB/IMG and reinstall.

Also, note that changing resolutions is supported but can sometimes lead to issues. If this is the case, reboot the server in Video Safe Mode.

UEFI Issues

The Arista UEFI Installer does not support UEFI SecureBoot, so SecureBoot must be disabled in your hardware's firmware menu prior to installing.

Unfortunately, not all UEFI implementations are written equally, which may cause issues when installing NG Firewall via UEFI on some hardware. If NG Firewall fails to successfully install via UEFI, check if your hardware's firmware can be configured for legacy BIOS boot and attempt to install using the normal installer. Otherwise, check out the Debian project's UEFI page for some tips on troubleshooting UEFI based installs.

1.3   Network Configuration

The most critical configuration in NG Firewall is the proper configuration of your network settings in **Config** > **Network**.

For simple networks, the configuration completed during the Setup Wizard is probably sufficient. However, some networks have multiple WANs, multiple LANs, various subnets, VLANs, VRRP, etc. This describes how networking operates and is configured in NG Firewall.

This section includes the following sections:

- Interfaces

- Hostname
- Services
- Port Forward Rules
- NAT Rules
- Bypass Rules
- Filter Rules
- Routes
- DNS Server
- DHCP Server
- Advanced
- Network Reports
- Troubleshooting

1.3.1 Cardinal Rules

There are several key rules to how NG Firewall operates that should be understood before deploying NG Firewall in an advanced/complex network.

NG Firewall MUST be installed in-line. NG Firewall is a gateway product, and it is designed to be in-line with network traffic. Some network administrators want to deploy some of the functionality of NG Firewall without installing it in-line. This isn't how NG Firewall is designed and it will likely not work.

For example, Spam Blocker will filter SMTP as it passes through NG Firewall. It will not store-and-forward to your email server like some products. Web Filter will filter web traffic as it passes through NG Firewall. It is does not operate as a explicit proxy that you "point" clients' browsers to send web traffic. All applications and functionality are designed to operate in a context where NG Firewall is installed in-line with the network flow of traffic.

NG Firewall MUST have a working internet connection. NG Firewall and many of its apps rely on cloud services. NG Firewall must have a working and consistent connection with the internet. This includes unfiltered HTTPS, HTTP, and DNS access to various cloud services. Without a valid internet connection and configuration many functions of NG Firewall will not work properly.

NG Firewall routes ALL traffic according to its routing table. Obviously, this is how all routers operate. They receive packets on an interface and then lookup in the routing table/rules where to send it. Where NG Firewall differs is that it is often installed as a bridge or with some interfaces bridged together. In the NG Firewall context, two bridged interfaces share a configuration (some products call them "zones"). Traffic passing between bridged interfaces are still subject to this cardinal rule.

This is often a surprise to people on complex networks as effectively you will need to tell NG Firewall where to send all the traffic on your network if you want it to go to the correct place. If you have a subnet that NG Firewall doesn't have a route for, then it will be sent to the default gateway even if that subnet is internal. For NG Firewall to operate correctly, you must configure it with a complete routing table so it knows how to reach all hosts on your network.

1.3.2 Placing NG Firewall on the Network

The first step, after understanding the above cardinal rules, is to decide where to place NG Firewall on the network. Obviously, NG Firewall must be installed in-line with all network traffic so this provides two options:

1. Install NG Firewall as the gateway/firewall for the network.
2. Install NG Firewall *behind* an existing gateway/firewall in flow with traffic.

Installing NG Firewall as the gateway/firewall is recommended. It is usually the simplest approach and it allows NG Firewall to leverage its full feature set including WAN Failover and WAN Balancer. This also places it in a convenient place to handle other separate internal networks (like wireless segments) that may only be connected at the gateway. Also, If you have tagged VLANs it is much simpler to run NG Firewall as a endpoint of those VLANs. This is commonly referred to as *router mode.*

However, often organizations don't want to replace the existing gateway/firewall or can't because it is controlled by a different organization. In these cases, installing NG Firewall as a "bridge" behind the gateway allows NG Firewall and the apps to scan and process network traffic without providing the routing functions of your firewall. This is commonly referred to as *bridge mode.*

Also, technically you can stick NG Firewall *in front of* an existing firewall. Typically firewall/gateways use NAT so all internal hosts share external IPs. This means by the time the traffic reaches the NG Firewall outside the firewall the source address of all internal communication will have the public address of the firewall. As such, NG Firewall can not differentiate between internal hosts so much of the functionality of NG Firewall (Web Filter, Reports, Shield, and so on) is severely compromised. Installing NG Firewall outside a NAT device is never recommended.

1.3.3 Configuring the Interfaces

The second major step, after choosing a place to deploy NG Firewall, is configuring the interfaces. All the configuration options of interfaces are documented in the Interfaces documentation. Likely, external and internal and are already configured from the Setup Wizard.

After the setup wizard you might still need to do some of the following configuration:

- Configure additional subnets on the external/internal.

Because NG Firewall routes all traffic according to its routing table, additional routes/aliases may be required for any additional subnets on your network. More information can be found on that in NG Firewall Installation.

- Configure additional interfaces

After the setup wizard only external and internal are configured. Additional interfaces are disabled and still require configuration. More information can be found on that in NG Firewall Installation.

- Configure any tagged VLANs

If you have tagged VLANs (802.1q) on your network, you will need to add VLAN Tagged Interfaces.

1.3.4 Bridging

When two interfaces are bridged in NG Firewall this means they are effectively sharing a configuration. Some products use the concept of "zones." In this terminology, bridging two interfaces puts those interfaces in the same "zone" or "network space."

Standard Bridge Mode

The most common scenario this is used is in *bridge mode* where the External is bridged to the Internal. This is extremely useful when there is a firewall upstream.

For example, if the firewall is *192.168.1.1*, then you can configure External as *192.168.1.2/24* with *192.168.1.1* as the gateway. The internal hosts all have *192.168.1.** addresses and can continue to use *192.168.1.1* as a gateway (*192.168.1.2* will also work as a gateway).

Gateway: 192.168.1.1

Address:
192.168.1.2/24

WAN

bridged

Internal

Internal Hosts
Address: 192.168.1.*/24
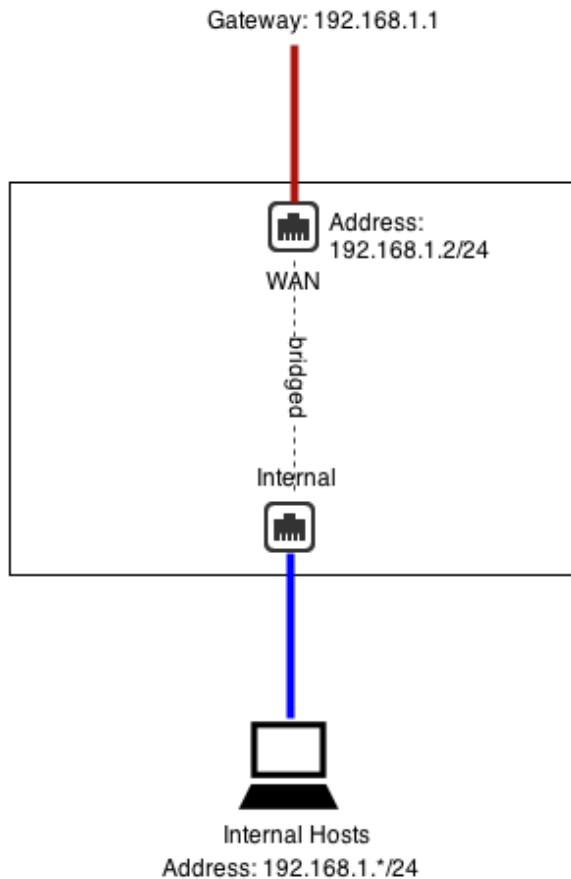
It is important to remember that even when bridging *NG Firewall routes ALL traffic according to its routing table.* This means if you have other subnets besides **192.168.1.*** like **192.168.2.***, then you will need to add aliases or routes for them otherwise that traffic will go to the default gateway.

DMZ Bridge

Another common scenario to use bridging is when NG Firewall has a public IP (**1.2.3.2** in this example), but you have other public servers with public IPs (**1.2.3.***). You could put those servers on the private network and use Port Forward Rules. But lets assume you wanted to keep them configured with public IPs to keep them separate from the internal and avoid any NAT/port forwarding issues.

In this case, you can bridge a "DMZ" interface to your external and it essentially shares the configuration and "zone" with external. This means you can place servers with public IPs on that segment and they can continue to use **1.2.3.1** as a gateway.

Gateway: 1.2.3.1

Address: 1.2.3.2/24

WAN

bridged

Internal

Address:
192.168.1.1/24

DMZ

Internal Hosts
Address: 192.168.1.*/24

Public Servers
Address: 1.2.3.*/24
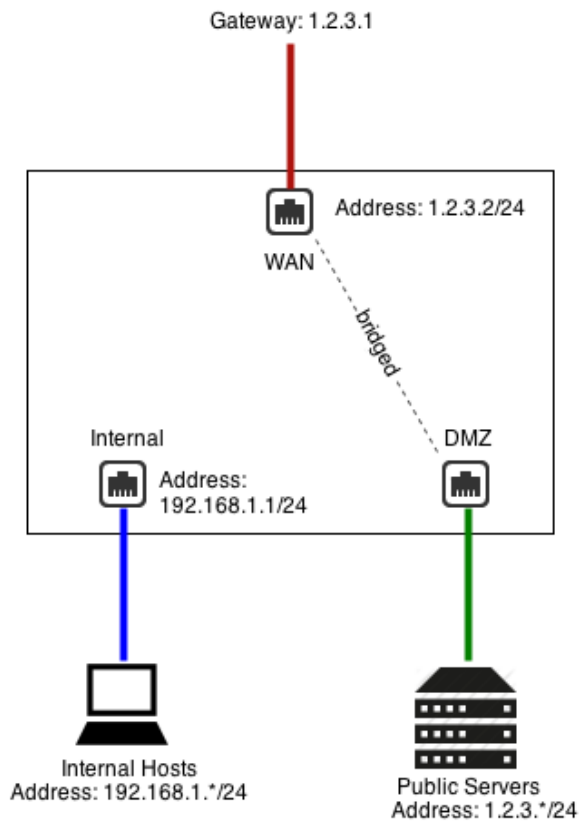
Additional Port

You can also just use bridge mode to provide alternate ports to existing interfaces/zones. Be careful, as traffic between the two goes through NG Firewall!

For example, if *Interface 2* is configured as *192.168.1.1/24* and *Interface 3* is bridged to *Interface 2*, then they are both effectively *192.168.1.1*. Basically *Interface 3* becomes an additional port for the *Interface 2* network.
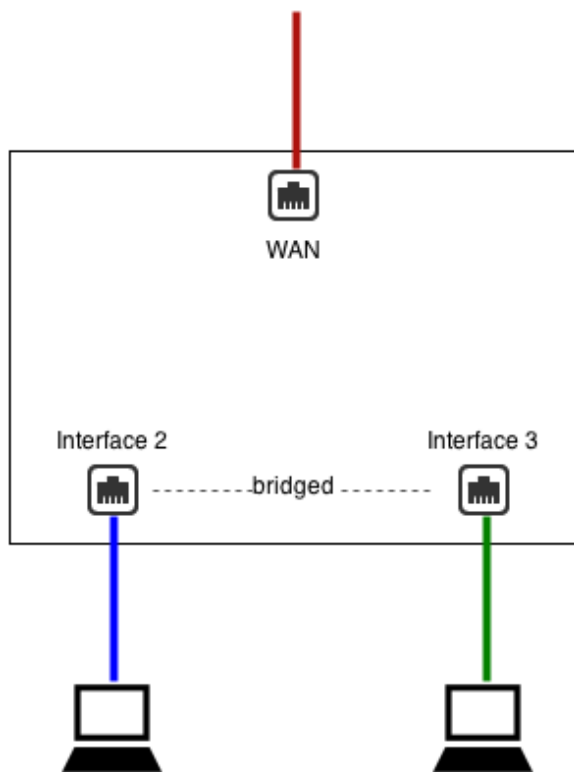
This is almost identical to a configuration without *Interface 3* where *Interface 2* is plugged into a switch with two free ports.
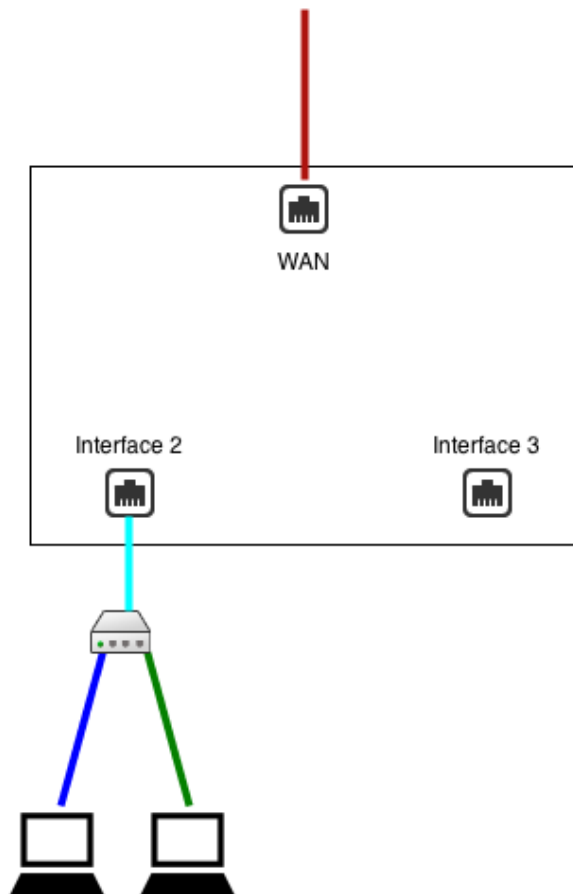
There are some important differences:

- In scenario 1, traffic between *Interface 2* and *Interface 3* goes through NG Firewall and is routed via NG Firewall's routing table
- In scenario 1, traffic between *Interface 2* and *Interface 3* goes through NG Firewall and is scanned by the apps (if not bypassed via Bypass Rules)

What "bridged" really means.

In NG Firewall when two interfaces are bridged it means that they are in the same zone or that they both connect to the same network space. As the cardinal rules explain, **NG Firewall routes all traffic according to its routing table** - even those crossing between two bridged interfaces. This is sometimes called *brouting* or a *brouter* - unlike how a traditional layer-2 bridge/switch behaves.

It means that packets coming inside one side of a bridge will NOT necessarily exit the other side of the bridge. It also means that packets destined with a specific route will be routed according to NG Firewall's routing table. **All traffic is routed according to NG Firewall's routing table.** It also means MAC addresses are not maintained across segments, even if they are bridged together as packets are brouted.

This may cause you to wonder how NG Firewall works in the traditional "Bridge Mode." The answer is simple, for outbound traffic to the internet NG Firewall will route that to its default gateway which was probably where the traffic was headed anyway and definitely where it should go. For inbound traffic NG Firewall knows where each local host on the bridged segment lives and it routes it directly. So inbound and outbound traffic both flow as expected.

Where things get complicated is when networks have more complicated routes and do not configure NG Firewall with those routes. Assume NG Firewall is installed in traditional bridge mode on a

192.168.1.1/24 network. Lets assume the network also has another internal network of *192.168.2.1/24* behind an internal router *192.168.1.100*. There is probably already a route on the existing firewall telling it that *192.168.2.** can be reached behind *192.168.1.100*. If the user then inserts NG Firewall in bridge mode and configures it as *192.168.1.2*, the entire *192.168.1.** network will work but none of the traffic on *192.168.2.** will work. Why? Because NG Firewall routes *all* traffic according to its routing table. The firewall will route *192.168.2.** traffic to *192.168.1.100*. When that traffic passes through NG Firewall it will not route it to *192.168.1.100* - it will route it according to its routing table. Since it knows nothing about *192.168.2.** and those addresses aren't local, it will be sent back out to the default gateway. As such the *192.168.2.** network will be completely offline as return traffic from the internet will not reach those hosts. Once a *192.168.2.0/24* route to *192.168.1.100* is added to NG Firewall traffic will flow as expected. The routing table on NG Firewall must reflect the layout of the network.

Another common scenario is bridging two separate networks with one NG Firewall server. Lets look at an example with 4 interfaces: *network1External*, *network1Internal*, *network2External*, *network2Internal*. *Network1Internal* is bridged to *network1External*. *Network2Internal* is bridged to *network2External*. The problem with this scenario is that **NG Firewall routes all traffic according to its routing table.** If traffic comes in on *network2Internal* and is bound for the internet it will **NOT** be sent out *network2External* just because that's where it was originally headed. It will be routed according to NG Firewall's routing table, which is the default route of NG Firewall - probably *network1External's* gateway! To setup this scenario one must use WAN Balancer and routes to assure that traffic coming in *network2Internal* is routed via *network2External*. This is true whether or not the separate network or separate physical networks or separate VLAN networks.

The key to using bridge mode effectively is understanding how NG Firewall routes traffic. While bridging can often be convenient it can also create headaches for complicated setups.

1.3.5   NAT

NAT or Network Address Translation is the operation of rewriting the source address of packets. Typically this is used so many internal hosts with internal IPs (*192.168.*.**, *10.*.*.**, and so on) can share one or several public IPs.

There are three ways that NAT is done is NG Firewall:

1. If you check *NAT traffic exiting this interface (and bridged peers)* on any WAN interface configuration.
2. If you check *NAT traffic coming from this interface (and bridged peers)* on any non-WAN interface configuration.
3. If you add a NAT Rule
4. NAT Rules

NAT traffic exiting this interface (and bridged peers)

The first option is *NAT traffic exiting this interface (and bridged peers)*. As described in the Interfaces documentation, this option will NAT any traffic exiting this interface and any of its bridged peers. This is enabled by default on WANs.

What this means is that any and all traffic exiting that WAN interface or bridged peers will be NATd to *auto* which is the current primary address of that WAN interface. Traffic between this interface and any bridged peers will not be NATd. Checking this option also blocks all traffic coming to this WAN that is not to a local process or explicitly forwarded with a Port Forward Rules.

NAT traffic coming from this interface (and bridged peers)

The second option is *NAT traffic coming from this interface (and bridged peers)*, as described in the Interfaces documentation, will NAT any traffic coming from this interface and any of its bridged peers. This is not enabled by default.

What this means is that all traffic from this interfaces will get NATd to *auto* which is the primary address of which ever interface the traffic it exits. Traffic between this interface and any bridged peers will not be NATd. Checking this option also blocks all traffic to this non-WAN except traffic forwarded with a Port Forward Rules.

When and Where to perform NAT

If there are only two interfaces in NG Firewall the first and second options are identical. When there are multiple internal subnets you will need to configure where you wish to NAT so that you have the desired behavior.

If you wish internal networks to be able to speak with each other (*i.e. **192.168.1.100** should be able to reach **192.168.2.100** on a different interface) then you do not want to NAT between those networks.
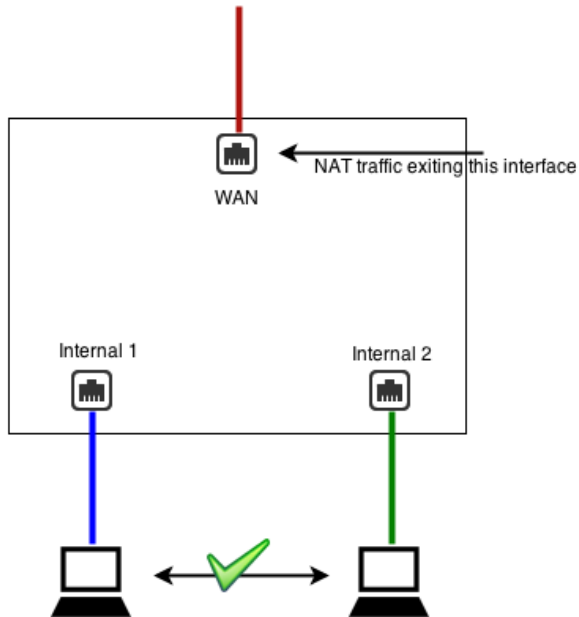
As such, you should uncheck *NAT traffic coming from this interface (and bridged peers)* on both LAN interfaces, and check *NAT traffic exiting this interface (and bridged peers)* on the WAN(s).

If you wish for internal networks to be completely separate such that they can not speak to each other then you want to check *NAT traffic coming from this interface (and bridged peers)* on the non-WAN interfaces. This means NAT will be performed on all traffic from these LANs and inbound sessions will be blocked unless explicitly forwarded with Port Forward Rules.
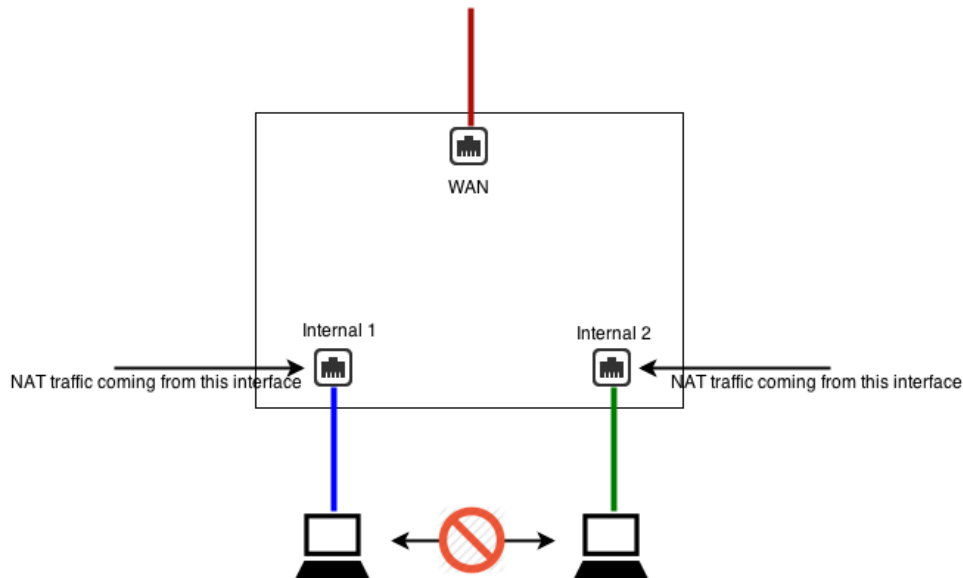
NAT Rules

The third option is explicitly configuring exactly what should be NATd to what address with NAT Rules. Be aware, these rules do not also explicitly block any inbound traffic like the two NAT options. NAT rules can be used in conjunction with the two NAT checkboxes and any matched NAT rule will take precedence. Most networks need only one the first two options, but sometimes there are scenarios where a NAT rule is wanted desired such as 1:1 NAT or when you want to guarantee that certain traffic (like SMTP exiting an email server) uses another address other than the primary WAN address.

If no NAT option is enabled (all unchecked and no NAT rules) then NG Firewall will route like a typical router without performing any NAT operation.

> **Note:** The *NAT traffic exiting this interface (and bridged peers)* option in WAN is equivalent to appending an *Auto* NAT rule to the end of the NAT Rules matching all traffic with *Destination Interface* equal to that WAN or any bridged peer but excluding traffic between any bridged peers in that zone. It also includes a Filter Rules to block inbound sessions from this WAN or bridged peers not explicitly port forwarded excluding sessions between any bridged peers in that zone.

> **Note:** The *NAT traffic coming from this interface (and bridged peers)* option in WAN is equivalent to appending an *Auto* NAT rule to the end of the NAT Rules matching all traffic with *Source Interface* equal to that non-WAN or any bridged peer but excluding traffic between any bridged peers in that zone. It also includes a Filter Rules to block inbound sessions to this non-WAN or bridged peers not explicitly port forwarded, excluding sessions between any bridged peers in that zone.

1.3.6 VLANs

VLANs or Virtual LANs are commonly used so that multiple subnets can share the same wire while maintaining complete separation including separate broadcast domains.

> **Important:** The term VLAN is sometimes also used to describe putting multiple untagged (no 802.1q tag) subnets on the same wire. For example, NG Firewall is in bridge mode as **192.168.1.2/24** but there is also a **192.168.2.*** on the same wire. If there are no 802.1q tags on the **192.168.2.*** traffic - it is **NOT** a VLAN and new VLAN interfaces should **NOT** be created on NG Firewall. In this scenario you should use guidance in NG Firewall Installation to properly configure NG Firewall to handle these subnets. VLAN interfaces will **ONLY** handle tagged 802.1q packets and all packets sent to a VLAN interface will be tagged with 802.1q tags.

VLANs have several uses. Often you want multiple completely separate internal subnets on a network, but you do not want to run multiple physical ethernet networks through a building. VLANs allow you to run multiple networks on the same physical wire while still guaranteeing they are completely separate and secure. To do this you must have VLAN enabled switches and products through-out the network. VLANs can also be useful if you have limited ethernet ports on NG Firewall and wish to overload a single NIC with two separate purposes. This requires that NIC to be connected to a VLAN enabled switch.

To create a VLAN interface click on the **Add Tagged VLAN Interface** button at the bottom of the interfaces grid on the **Config** > **Network** > **Interfaces**. This will create a new *virtual interface.* First, you will need to give this interface a new name, and also select the **Parent Interface.** The **Parent Interface** is the physical interface that this VLAN virtual interface exists on. Then you will need to configure an *802.1q Tag* which is an integer between *1* and *4094* inclusive. After this you will configure the interface just like any interface.

This new *VLAN interface* is just like a physical interface in all ways. This means you can configured this VLAN interface exactly like any physical interface. It is completely separate from the physical parent interface. Any packet coming in on the physical parent interface with the 802.1q tag matching the configured value (*1-4094*) will be considered by NG Firewall to be coming in the VLAN interface. Any packet sent to the virtual VLAN interface will actually be sent on the physical parent interface with the configured 802.1q tag. All untagged packets on the physical parent interface will be processed like normal through the physical parent interface, only 802.1q tagged packets with the matching 802.1q tag will be processed by the VLAN interface.

After configuring NG Firewall with the appropriate tagged VLAN interfaces, you will need to configured some VLAN-enabled managed switch to properly process the packets as desired. For example, how to configure a port based vlan on a hp procurve, describing how to configure a HP ProCurve switch.

> **Note:** VLAN interfaces are completely separate from their physical parents, however they do share the same physical NIC and as such will be limited by the throughput of the physical NIC. For example, if you have two 100Mbit tagged VLAN WANs on the same physical 100Mbit NIC, then you will still be limited to 100Mbit total on both WANs at any given time.

Example: We want to have two completely separate LANs on our network but we only have one wire or one network card. As such we need a VLAN enabled switch. Configure the internal interface with the IP and configuration from *LAN 1*. Create a tagged VLAN interface with 802.1q tag *3.* Configured the new VLAN interface with the IP and configuration for *LAN 2*. Now configure your VLAN switch to send *VLAN 3* packets to the appropriate ports with *LAN 2* hosts. In this scenario we are using the same wire but have two separate LANs with separate broadcasts domains.
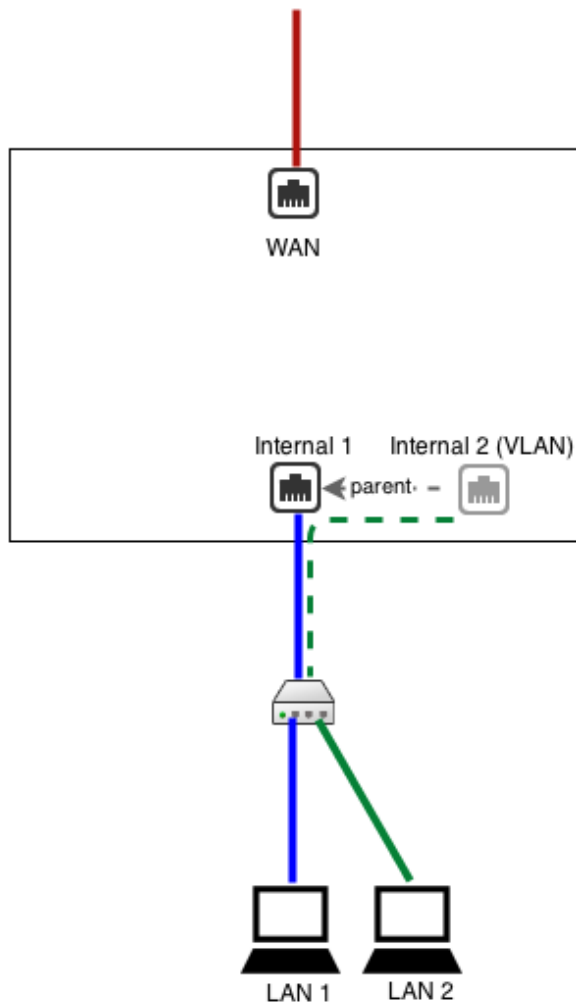
> **Note:** If you want to use a single wire and/or network card but you don't care about keeping the two LANs seperate, you don't need VLANs and can just use aliases/routes.

Configuring VLAN on NG Firewall in Bridge Mode

Some users want to configure NG Firewall in bridge mode in the middle of a VLANed network. This is possible, but **NOT RECOMMENDED**. It is suggested to install NG Firewall as the gateway and terminate VLANs on an addressed VLAN interface. However, if you wish to install NG Firewall as a bridge in the middle of several (V)LANs simultaneously, the following instructions will allow you to establish multiple bridges and then enter the routes to tell NG Firewall that if traffic enters on a specific interface it should exit the bridged peer.

• You will need to create two virtual interfaces for each VLAN you want to set up.

  1. One as a child to the external interface
  2. One as a child to the internal interface
• To set up the external's VLAN interface

  1. Click on **Create a tagged VLAN interface**
  2. Give the interface a name that's easily identifiable by you
  3. Set the Parent Interface to External
  4. Set the 802.1q tag
  5. Config type must be "Addressed"
  6. Under IPv4 configuration, assign a unique static IP to the interface.

7. Enter the IP address for the VLAN gateway
8. Enter DNS servers you would like to use

- To set up the internal's VLAN interface

  1. Click on **Create a tagged VLAN interface**
  2. Give the interface a name that's easily identifiable by you
  3. Set the Parent Interface to Internal
  4. Set the same 802.1q tag that you configured on the external VLAN interface
  5. Config type must be "Bridged"
  6. Bridge this interface to the external VLAN interface

- Go to WAN Balancer and set the weights to send 100% of network traffic out of your untagged external interface.
- Click over to the **Route Rules** tab and create a new rule.

  1. Source Interface : [Internal VLAN interface]

2. Destination WAN : [External VLAN interface]

1.3.7 VRRP

VRRP provides network level redundancy.

Multiple NG Firewalls can be run in parallel in a high availability configuration. In this configuration one NG Firewall will be the "primary" and one or more NG Firewalls will be the "secondary." In the event the primary fails, one of the remaining secondary NG Firewall will take over the primary role such that network traffic continues to flow without interruption.

NG Firewall uses Virtual Redundancy Router Protocol or VRRP to handle the switching between NG Firewall servers. All NG Firewall servers must be on and all configured with a share VRRP Virtual Address. The primary is the only NG Firewall to answer/handle traffic routed to the VRRP Virtual Address. If the primary fails an "election" is held over VRRP and the next-highest-priority secondary will begin handling traffic to the VRRP Virtual Address.

All NG Firewall interfaces must be configured statically, and there must be no bridged interfaces. Parallel NG Firewalls configured as bridges will create a bridge loop!

VRRP Basic Example

A common configuration is running two NG Firewalls to act as the gateway for the internal network. For example, let's assume *NG Firewall 1*, the primary, has a public IP of *1.2.3.4* and an internal IP of *192.168.1.2*. Lets assume NG Firewall 2 has a public IP of *1.2.3.5* and an internal IP of *192.168.1.3*. Both are running in "router" mode and doing NAT and acting as a gateway. Note that all IPs must be unique! This configuration requires each NG Firewall to have its own external IP!

Now we configure *NG Firewall 1* to have a VRRP Virtual Address of *192.168.1.1* on the Internal interface, and also configure NG Firewall 2 to also have a VRRP Virtual Address of *192.168.1.1* its Internal interface. They both share the same VRRP Virtual Address. Each NG Firewall in the group must have the same VRRP ID. So let's give *NG Firewall 1* a VRRP ID of *1* and also give *NG Firewall 2* a VRRP ID of *1*. We want *NG Firewall 1* to be the primary when it's on and working without issues so give it a higher priority of *100*. *NG Firewall 2* is the secondary so it should be given a lower priority: let's give it a lower priority of *50*.

Internet

1.2.3.4

1

192.168.1.2

VRRP Virtual Address
192.168.1.1

1.2.3.5

2

192.168.1.3

Internal Network

☑ Enable VRRP

VRRP ID: 1

VRRP Priority: 100

VRRP Aliases

Add		Import Export
Address	Netmask / Prefix	Delete
192.168.1.1	24	✕

☑ Enable VRRP

VRRP ID: 2

VRRP Priority: 50

VRRP Aliases

⊕ Add		🡄 Import	🡆 Export
Address	**Netmask / Prefix**		**Delete**
1.2.3.3	24		✕

Configure your internal hosts to use the VRRP Virtual Address (*192.168.1.1*) as the gateway. In this configuration the primary will route all traffic to *192.168.1.1* just like a regular address. Should the primary fail within a few seconds the secondary will become the new primary and start routing traffic.

Note: You should configure the DHCP server to hand out *192.168.1.1* as the default gateway. If NG Firewall is providing DHCP. Configure *NG Firewall 1* as the "authoritative" with *192.168.1.1* as the "Gateway Override." Configure *NG Firewall 2* the same but as non-authoritative. This way *NG Firewall 1* will handle all DHCP unless it is down, in which case *NG Firewall 2* will handle DHCP.

VRRP External Example

The above example work great for outbound traffic, but if you have inbound traffic being port forwarded that traffic might fail if the NG Firewall owning that address fails. VRRP can also be used to provide redundancy for inbound traffic. For example, similar to above, let's assume you have *NG Firewall 1* with *1.2.3.4* and *NG Firewall 2* with *1.2.3.5*.

You can configure both with a shared VRRP Virtual Address of *1.2.3.3*. Now configure port forwards on both NG Firewall for traffic destined to *1.2.3.3* to the appropriate internal host. Only the primary will handle traffic to *1.2.3.3*. If the primary fails the secondary will take over traffic handling for *1.2.3.3* and port forwarding. External hosts will still be able to reach local services should the primary fail.

In this scenario NAT Rules can also be configured if outbound traffic should use the same address regardless of which server is handling the traffic.

VRRP Combined Example

It is also possible to combine VRRP on multiple interfaces. For example, combine the above two examples. VRRP can be used to provide redundancy on both interfaces. VRRP IDs must be unique for each server. For example, the external on both should be *VRRP ID 1* and the internal on both should be *VRRP ID 2*. In this scenario VRRP is "grouped" such that if the server loses its "primary" status on one interface it will also release its primary status on other interfaces.

Internet

VRRP Virtual Address
1.2.3.3

1.2.3.4

1.2.3.5

1

2

VRRP Virtual Address
192.168.1.1

192.168.1.2

192.168.1.3

Internal Network

☑ Enable VRRP

VRRP ID: 2

VRRP Priority: 100

VRRP Aliases

Add		Import	Export
Address	Netmask / Prefix		Delete
1.2.3.3	24		✕

☑ Enable VRRP

VRRP ID: 2

VRRP Priority: 50

VRRP Aliases

⊕ Add	Import Export
Address	**Netmask / Prefix**	**Delete**
1.2.3.3	24	✕

☑ Enable VRRP

VRRP ID: 1

VRRP Priority: 100

VRRP Aliases

⊕ Add	Import Export
Address	**Netmask / Prefix**	**Delete**
192.168.1.1	24	✕

☑ Enable VRRP

VRRP ID: 1

VRRP Priority: 50

VRRP Aliases

⊕ Add	Import Export
Address	**Netmask / Prefix**	**Delete**
192.168.1.1	24	✕

For example, in the picture above if the Internal interface on *NG Firewall 1* is unplugged, then *NG Firewall 2* will become the primary and start responding to *192.168.1.1*. *NG Firewall 1* will also release its primary status on the external interface so that *NG Firewall 2* will also handle *1.2.3.3*. This is to avoid any scenarios where *NG Firewall 1* is primary the external address and *NG Firewall 2* is primary on the internal address.

1.4   NG Firewall User Guide

NG Firewall User Guide

Contents

- Administration Interface

    The Administration Interface is the main interface used to configure NG Firewall.

- Dashboard

    The *Dashboard* provides an overview of the state of your NG Firewall. It is extremely useful for quickly viewing or monitoring what is happening on the network and the current status of the NG Firewall server.

- Applications

    The Administration Interface is the main interface used to configure NG Firewall.

- Config

    The config tab holds all the settings related to configuration of the NG Firewall server itself and settings for components of the *platform* that apps may interact with.

- Reports

    Reports provides users with detailed statistics of the traffic and activity on your network.

1.4.1   Administration Interface

The Administration Interface is the main interface used to configure NG Firewall.

A registration and welcome message is displayed upon the first visit to the administration interface. It will make a suggestion about the suggested applications that may be useful for your network. You can choose to install the recommended apps or to install the apps manually.

There are four main tabs in the administration interface in the main menu:

1. Dashboard
2. Apps
3. Config
4. Reports (only visible if the Reports app is installed.)

In the sub-menu there are four views:

1. Sessions

2. Hosts
3. Devices
4. Users

Tip: Using Mozilla Firefox or Google Chrome browsers is recommended for administration.

1.4.2   Dashboard

The *Dashboard* provides an overview of the state of your NG Firewall. It is extremely useful for quickly viewing or monitoring what is happening on the network and the current status of the NG Firewall server.
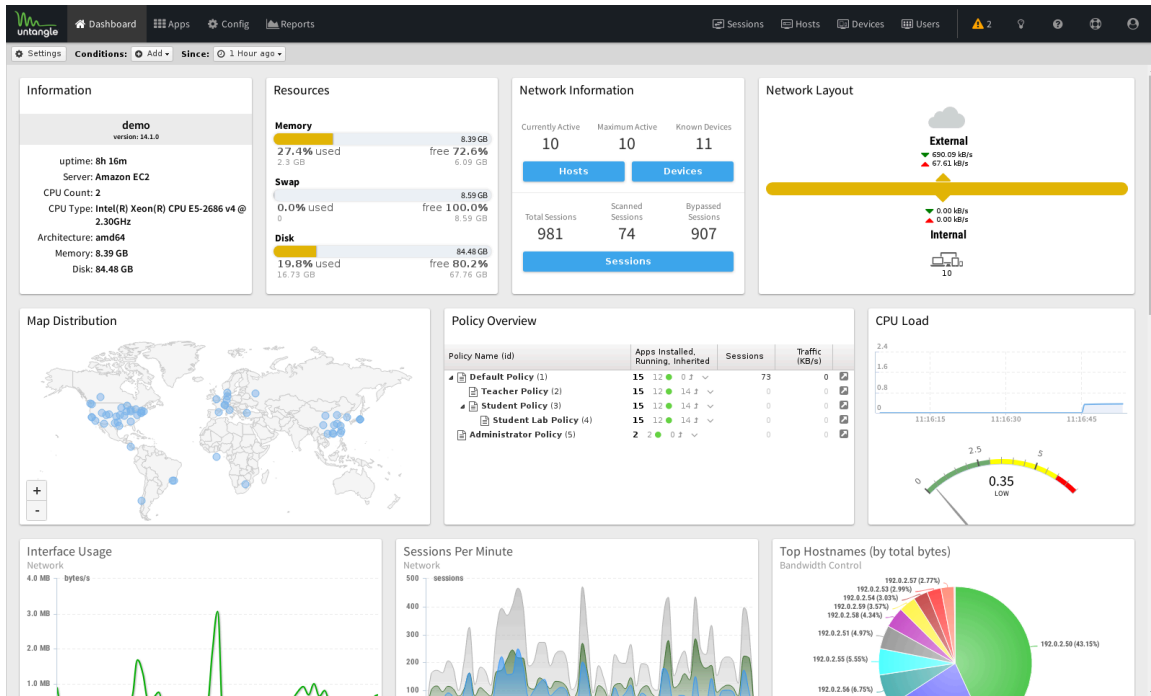
Figure 11: NG Firewall Dashboard

By default the dashboard will show several *widgets* with varying pieces of information. However, the dashboard is completely customizable. Widgets can be removed and added so the administrator sees exactly the information that is important to them on the dashboard.

There are many different type of *widgets* available:

Name	Information
Information	Shows some information about NG Firewall, like name, model, version, etc.
Resources	Shows an overview current memory swap and disk usage.
CPU Load	Shows a graph of recent CPU load.
Network Information	Shows an overview of the network information like session count and device/host count.
Network Layout	Shows an overview of the network layout based on the interface configuration.
Map Distribution	Shows the current sessions mapped geolocation on a world map, sized by throughput.

Name	Information
Report	Shows any Report Entry from Reports

To change what is displayed on the dashboard click on **Manage Widgets** at the top. From here you can show or hide the built-in widgets, or add new widgets from Reports by clicking on the **Add** button.

When adding a Report widget, you must also specify a timeframe (the number of hours worth of data to display) and a refresh interval (how often the widget refreshes on the dashboard).

When viewing a Report Entry in Reports you can also easily add it to your dashboard by clicking on the **Add to Dashboard** button.

If you see an alert icon near the top of the dashboard, hover over it for more information - these are Administrative Alerts designed to help you keep your NG Firewall healthy.

1.4.3   Applications

Applications are plugins that add functionality to your NG Firewall server - just like "apps" on an iPhone or Android device.

On the Apps tab you'll see the currently installed apps. Across the top there is a dropdown menu to switch to different **Policies.** Policies can be controlled via the Policy Manager app.

Apps can be installed by clicking the **Install Apps** button at the top. It will display the apps that can currently be installed. To install an app simply click on its icon. You can install as many apps at the same time as you like. After starting the installation of the desired apps you can click the **Done** button at the top to go back to the apps view.

Once installed the applications' settings can be configured by clicking on the **Settings** button or the icon of app, depending on the skin. Applications install with the suggested configuration which is the default settings and on/enabled in most cases. An application that is off/disabled will not process any

network traffic. To enable a disabled application edit the settings and click **Enable** on the first tab inside the settings.

After clicking **Settings**, you will be presented with tabs for different settings sections, as well as typical buttons marked **OK**, **Cancel** and **Apply**. **Apply** saves any changes. **OK** saves any changes and closes the window. **Cancel** closes the window without saving settings. On the left hand side there is a **Remove** button which will remove the application from the current policy. The **Help** button will open the help for the tab currently being viewed.

NG Firewall has two types of Applications:

• **Filter Applications** All the Applications *above* the **Services** pane in the interface can have one instance per policy.
• **Service Applications** All the Applications *below* the **Services** pane are global and exist in all virtual racks.

Many networks only need one *policy* which means all traffic gets processed by the same apps and same configuration, but multiple policies (sometimes called "racks") are possible for bigger networks. For more information about running multiple racks check out the Policy Manager application.

To learn more about each application use the links below.

Filter Applications

Web Filter	Web Monitor	Virus Blocker	Virus Blocker Lite	Spam Blocker	Spam Blocker Lite
Phish Blocker	Web Cache	Bandwidth Control	Application Control	Application Control Lite	SSL Inspector
Captive Portal	Firewall	Intrusion Prevention	Threat Prevention	Ad Blocker	

Service Applications

Reports	Policy Manager	Directory Connector	Web Monitor	WAN Balancer
Captive Portal	IPsec VPN	OpenVPN	WireGuard VPN	Branding Manager
Configuration Backup	Live Support			

1.4.4   Config

The config tab holds all the settings related to configuration of the NG Firewall server itself and settings for components of the *platform* that apps may interact with.
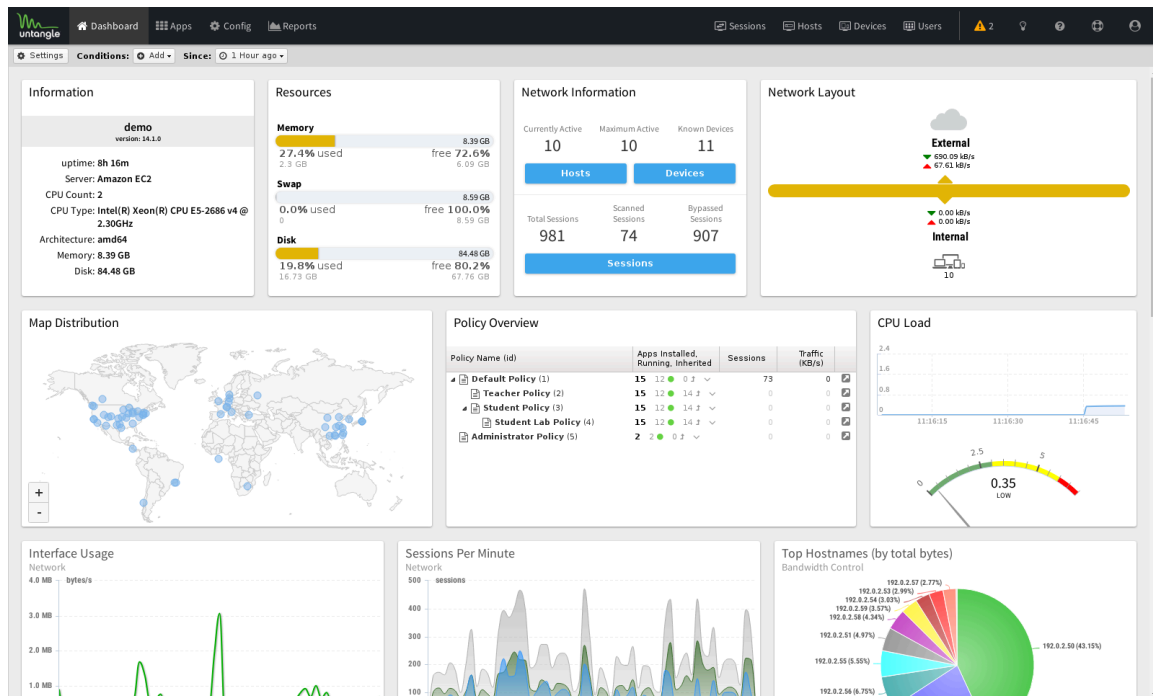
This is an list of all sections available under the **Config** tab in the Administration UI.

Network

The *Network* configuration contains all the settings to control how your NG Firewall server routes and handles network traffic. Properly configuring network settings is critical for proper operation.

- Interfaces
- Hostname
- Services
- Port Forward Rules
- NAT Rules
- Bypass Rules
- Filter Rules
- Routes
- DNS Server
- DHCP Server
- Advanced

  - Options
  - QoS
  - Access Rules
  - UPnP
  - Network Cards
  - DNS and DHCP
  - Netflow
  - Dynamic_Routing
- Network Reports
- Troubleshooting

The Network Configuration documentation documents how networking in NG Firewall functions and is commonly configured.

Administration

Administration controls the administration-related functionality of the NG Firewall server.

- Admin
- Certificates
- SNMP
- Skins
- Google

Email

Email contains all the email-related configuration of the NG Firewall server.

- Outgoing Server
- Safe List
- Quarantine

Local Directory

Local Directory stores a list of users that can be used by the applications. It also supports RADIUS for 802.1x authentication from properly configured wireless network access points.

The RADIUS Server can be enabled to allow WiFi users to authenticate as any user configured in *Local Directory*.

The RADIUS Proxy can be enabled to allow WiFi users to authenticate with credentials that are validated with a configured Active Directory Server.

- Local Users
- RADIUS Server
- RADIUS Proxy
- RADIUS Log

Upgrade

Upgrade allows the server to upgrade and contains upgrade-related settings.

Upgrade Settings

Upgrades show the currently available upgrades if any. If upgrades are available, an upgrade can be started by pressing the *Upgrade* button at the top under Status.

To see changes see the Changelogs.

After the upgrade begins, it will download the new packages (which may take some time) and then the upgrades will be applied. Do not reboot or power off the server during the upgrade.

If *Automatically Install Upgrades* is checked, NG Firewall will automatically check for new versions and upgrade if available.

Automatic Upgrade Schedule configures when NG Firewall will automatically upgrade if upgrades are available. NG Firewall will automatically upgrade at the specified time on the days of the week than are checked.

Upgrade FAQs?

When will I get the upgrade version?

- Upgrades are rolled out over time to NG Firewall deployments. The rollout can stretch out several weeks. If you want the upgrade immediately, email your UID to the Support team and request they add your UID to the Early Upgrade list.

How do I know when the new version is available for my NG Firewall?

- In the Upgrade page of your NG Firewall, when a new version is available, the Upgrade button will appear. If automatic upgrade setting is enabled, your NG Firewall will upgrade automatically once the upgrade is available on day and time specified.

Does the upgrade require a reboot?

- If a reboot is needed, the upgrade will reboot automatically once the upgrade is installed. There is no need for a manual reboot. Most upgrades will not reboot as there is no kernel change.

How long does the upgrade take?

- It's difficult to be precise since customer platforms, Internet connection speed and complexity of the upgrade vary. Generally upgrades take less than 20 mins. If the database version is changed as part of the NG Firewall upgrade, the process will take longer as the database will need to be converted. There are extreme cases of the upgrade taking over an hour.

Do I need to reinstall?

- No, upgrade process will update all the components on the NG Firewall seamlessly.

Where can I get what is changed in the new version?

- Release changes are posted on the NG_Firewall_Changelogs page.

System

System contains settings related to the server

- Regional
- Support
- Logs
- Backup
- Restore
- Protocols
- Shield
- System Reports

About

About contains system information

- Server
- Licenses
- License Agreement

Reports

The reports tab is only visible if the Reports app is currently installed. To read more about reports view the reports documentation.

1.4.5   Reports

Reports provides users with detailed statistics of the traffic and activity on your network.

Contents

- About NG Firewall Reports
- Settings

    - Status
    - All Reports
- Report Entry

    - Data
    - Email Templates
    - Reports Users
    - Name Map
- Accessing Reports
- Report Viewer
- Conditions

    - Condition Operators
    - Conditions Example - Policy by Policy ID
    - Conditions Example - Web Filter Categories
- Related Topics

About NG Firewall Reports

These reports can be viewed online, either through the administration interface or through the separate reporting interface available to non-administrators reporting-only users.

Customizable report summaries can be sent via email, which includes basic information and a link to view the online reports if the user has access.

Reports can backup your data in multiple formats to Google Drive for long term storage.

Settings

This section reviews the different settings and configuration options available for Reports.

Status

On this tab you can click **View Reports** to open up Reports in a new browser tab.

Figure 12: Reports Status tab

All Reports

This is the full list of all currently existing reports. This includes all the default reports and any custom reports that have been added.

To edit a report, click on **View** and then click on **Settings.**

To delete a custom report, click on **View** and then click on **Delete.**

To create a new custom report, click **View** on an existing similar report, then click **Settings.** then change the name and click **Save as New Report**.

To create a report from scratch, go to **Reports** and click on **Create New** in the lower left. When creating reports from scratch, each field must be carefully chosen and tuned until the desired data is provided. This process can be time consuming and difficult. It is suggested to work with a similar report to required the desired result. Additionally you can ask for help via support or the forums and import the report if someone can craft it for you.

If creating a report from scratch, the settings and fields and their purposes are described below.

Figure 13: All Reports tab

Report Entry

A report is many settings that essentially describe how to craft a SQL query and how to display the data. Here are the fields:

Name	Value	Available	Description
Report Type	Text, Pie Graph, Time Graph, Time Graph Dynamic, Event List	The type of graph	
Title	Text	All	The report title
Category	Any existing category/ application	All	The category in which the report is located
Description	Text	All	A brief description of the report
Text String	Text	Text	The text used to create the Text Report Type
Pie Group Column	Text	Pie Graph	The column to "group by" in top X charts (usually user, host, and so on)
Pie Sum Column	Text	Pie Graph	The column to sum in the top X charts (usually count, bytes, and so on)
Order By Column	Text	Pie Graph	The column to order by
Graph Style	Pie, Pie 3D, Donut, Donut 3D, Column, Column 3D	Pie Graph	The render style of the pie graph
Pie Slices Number	Integer	Pie Graph	The number of slices to display
Units	Text	Pie Graph	The units being displayed (usually bytes, sessions, and so on)
Graph Style	Line, Area, Stacked Area, Column, Overlapped Column, Stacked Columns	Time Graph	The render style of the time graph

Name	Value	Available	Description
Time Data Interval	Auto, Second, Minute, Hour, Day, Week, Month	Time Graph	The time aggregation unit or resolution
Approximation	Average, High, Low, Sum	Time Graph	The method used to aggregate/combine data points
Units	Text	Time Graph	The units being displayed (usually bytes, sessions, and so on)
Series Renderer	None, Interface, Protocol	Time Graph	The renderer used to display human-readable names
Dynamic Column	Text	Time Graph Dynamic	The column to select for/ group by
Dynamic Value	Text	Time Graph Dynamic	The value to sort by and display
Dynamic Limit	Integer	Time Graph Dynamic	The number of series to show
Aggregation Function	Count, Sum, Min, Max	Time Graph Dynamic	The function used to aggregate by dynamic values grouped by dynamic column
Graph Style	Line, Area, Stacked Area, Column, Overlapped Column, Stacked Columns	Time Graph Dynamic	The render style of the time graph
Approximation	Average, High, Low, Sum	Time Graph Dynamic	The method used to aggregate/combine data points
Units	Text	Time Graph Dynamic	The units being displayed (usually bytes, sessions, and so on)
Series Renderer	None, Interface, Protocol	Time Graph Dynamic	The renderer used to display human-readable names
Colors	Color Picker	All	The color palette to use
Display Order	Integer	All	The integer used to determine the report's position in the category list

Data

Data Retention: This value controls how much time report data is kept on disk. Please note that increasing the number increases the amount of disk space that is needed for data storage.

> **Note:** NG Firewall version *16.3* and above stops reporting data when free space falls below *5* GB.

- *Delete All Reports Data*: If you are running low on disk space and you wish to free space by wiping the reports database you can use this option.

Google Drive Backup: If your system is connected to your Google account, you can configure Reports backups to Google Drive.

- *Upload Data to Google Drive* If enabled, and the Google Connector in Directory Connector is enabled, your daily data will be uploaded to google drive each night for safe storage.
- *Upload CSVs to Google Drive* If enabled, and the Google Connector in Directory Connector is enabled, your daily CSV files will be uploaded to google drive each night for safe storage.

- *Google Drive Directory* configures which subdirectory data will be uploaded to in google drive.

Import/Restore Data Backup Files imports data from a previous backup into the database.

> **Note:** This directly imports the SQL contents. If you have upgraded and the database schema has significantly changed since the time of the back, the import will not work correctly.

Figure 14: Reports - Data tab

Email Templates

You can customize emailed reports using Report Templates. You can create as many as you wish with any combination of:

- **Interval**: Daily, Weekly, Monthly, Week to Date, Month to Date. The determines the time interval that the report will cover. Beware that enough data is available via the Retention settings to provide the data for the configured interval.
- **Mobile**: Generate chart images more appropriate for a mobile device.
- **Reports**: Select those reports under Config and Application sections. Text and chart reports are allowed but not event list reports. Reports for applications will be included only if that application is installed.

Additionally, you can copy the settings for an existing report.

The default Daily Reports template includes common text and chart reports for your system. This template is fixed and cannot be changed or modified.

Email Templates must be associated with Report Users.

Figure 15: Email Templates tab

Reports Users

Reports users are users that are not administrators, but can still view reports.

- **Email Address** is the email address (and username) of the report user. *Admin* is a special case that determines if administrators will receive emails and alerts.
- **Email Alerts** determines if this reports user will receive email alerts.
- **Email Reports** determines if this reports user will received email report summaries.
- **Email Templates** determines which email report summaries this user will receive if *Email Reports* is enabled.
- **Online Access** if enabled, a URL to online reports is included in emailed report summaries for this user.
- **Change Password** changes the password for this reports user.

Figure 16: Reports Users tab

Name Map

You can use the Name Map to manually configure the hostname for hosts. NG Firewall often can automatically determine the hostname for the IP automatically via DHCP or other methods. You can view the current names for currently active hosts in the Hosts.

However, when NG Firewall is unable to automatically determine a hostname for an IP the Name Map provides a way to manually name them.
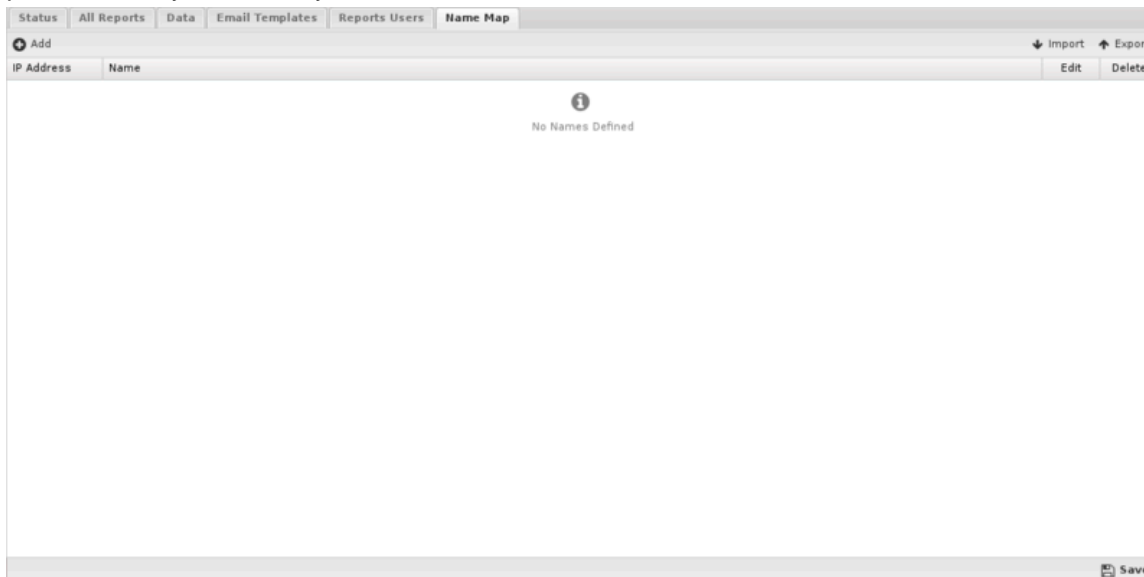
Figure 17: Name Map tab

Accessing Reports

If a user is set up to receive email report summaries, they only need to view or download the HTML attachment to see an overview report. If they need more information or would like to drill down to specific users or machines, they can use the link in the email, which will open Reports on the NG Firewall if it is accessible from their location.

To access Reports directly from a browser, you have two options:

- **Outside the NG Firewall network**: Browse to the IP of the NG Firewall /reports using HTTPs, such as https://1.2.3.4/reports.
- **Inside the NG Firewall network**: Browse to the IP of the NG Firewall /reports, such as https://192.168.1.1/reports.

Note that to view Reports from outside the network you'll need to check **Allow HTTPS on WANs** at **Config > Network > Advanced > Filter Rules**. If you have changed the **External HTTPS Port**, you'll need to use the proper HTTPS port when connecting from the outside.

Figure 18: Access Reports Summary

Report Viewer

Reports provide a graphical view of the network traffic and actions of your NG Firewall. Various reports are available within applications and base system components. The reports can be manipulated to drill down, customize, and export data in many ways using the Report Viewer.

There are a few panels in the Report Viewer:

- **The top panel**: This top panel (just below the navigation menu) allow you specify which data is viewed. By default, there is just a timeframe and no conditions, so reports will show all data for the specified timeframe. Conditions can be viewed to view more specific data, such as a specific host, user, domain, application, web category, etc.
- **The left panel**: This allows you to choose the report you wish to view. At the bottom you can use the search box to quickly find reports with the specified string in the title. You can also import and create new reports using the "Add/Import" button.
- **The chart panel**: This panel shows you the specified report. It also includes several action buttons at the top.
- **The data panel**: The data panel, hidden by default, can be displayed by clicking on the **Data View** button in the chart panel. This will show the raw data used to generate the chart and allow the user to export the data by clicking the **Export Data** button at the bottom.

Conditions

The Conditions panel appears at the top panel and can be used to filter data displayed in reports. For example, to view a "specific" host's report, you can add a condition for Client = "*192.168.1.100*" and then all reports available will only show data where the client is *192.168.1.100*. Multiple conditions can be added to drill down and inspect data. Conditions can also be added quickly by clicking on slices in pie charts.

The Add Condition dropdown contains many commonly used conditions, or the full list of all tables and columns can be browsed by clicking on the **More** button to add conditions for any database column.

> **Note:** Conditions will not apply to all reports. For example, If viewing a specific users report by adding a condition where *Username* = *foobar* - many reports will be greyed out and unviewable. This is because the data used to generate those reports is not relevent to the specific user (it

does not contain a username column). For example, the CPU usage report is a system report that is not relevant to a specific network user and so there is no way to filter that data by user.

Condition Operators

The second field in the condition is the logical operator that will be used in evaluating the condition value defined in the last field. In most use cases the default "=" operator is what you want to use. However, there are several other operators available that make the reports and alerts a whole lot more powerful.

Conditions Example - Policy by Policy ID

In many cases, you may just want to see the traffic related to a specific policy within Policy Manager. This can be accomplished very easily by adding a condition using the Quick Add feature.

1. In the Conditions panel, select **Add**.
2. Choose **Policy ID** and specify equals and the policy ID in question.
3. The conditions is applied and will remain applied as you switch between reports.

Conditions Example - Web Filter Categories

From pie charts, you can quickly add a condition from the Current Data window. This can be handy for use with the Web Filter category selection which we'll use for this example. Once the condition is applied, we can then use other reports to drill down to find out more information about the traffic such as which user might be responsible.

1. Open Report Viewer or the Web Filter Reports tab.
2. Select the **Top Categories** report (by size or requests). In our example, you can see Games was at the top.
3. Click on the Games pie slice, and when prompted to add a condition click **Yes**.
4. All Reports can now be viewed for Games only traffic.
5. For example, the Top Clients (by request) will show the clients that visited the most gaming sites.
6. For example, the Web Usage (scanned) will show "Gaming" web usage throughout the day of the network.

Related Topics

Custom Reports

1.5   NG Firewall Virtual Appliance on VMware

NG Firewall can be virtualized through a virtual appliance running on VMware ESX or ESXi.

The virtual appliance can also be used in for demonstrations in VMware player, workstation, fusion, or server, but it is not recommended run a production installation in these environments. Support will help with **NG Firewall** configuration but configuration of the virtualization hypervisor is beyond the scope of Edge Threat Management support.

- **Demo virtual appliance**: suitable for installation on a laptop or desktop in order to have a working instance of the platform running inside your Window, OS X, or Linux OS for testing or demonstration purposes. This is supported using VMware Player, Fusion, Server, or Workstation and requires only one physical network interface. Use this mode if you have only one physical network interface in your VMware host machine.
- **Production virtual appliance:** to be used as a network gateway. This mode requires at least two physical network interfaces (three if you want or need an external DMZ). We recommend you use either VMware ESX or ESXi Server. Use this mode if you have two or

more physical network interfaces that you can connect to external, internal and (optionally) DMZ networks.

NG Firewall Support and VMware

Arista Edge Threat Management wants you to have a successful deployment. Unfortunately, our support staff doesn't have the expertise in VMware ESX to ensure that we can help you with installing and configuring VMware. We will certainly help you with your NG Firewall configuration, provided it's running on ESX.

That being said, we'd like to make you aware that systems like NG Firewall that require a lot of real time processing aren't great candidates for virtualization. VMware works by "time-slicing" the physical CPUs in the host system. While the VMware server is off processing other virtual machines, the NG Firewall server is unable to process traffic. At the same time, network traffic continues to arrive. This traffic stacks up and presents itself to the NG Firewall VM as "bursty." This exacerbates any high load issues that may be present. The exact threshold of where it will be unsuitable is hard to say. It is a combination of traffic level, types of traffic, and user expectations.

In summary: We do not recommend virtualizing NG Firewall. If you choose to install NG Firewall in a virtual environment, the support team will assist you with any issues related to the NG Firewall and its applications, but they will not help with virtualization set up/connectivity issues or issues caused by virtualization (high load, slow speeds, etc).

How to Install on ESX or ESXi

Before you get Started

Requirements:

1. ***VMware ESX server version 6.5.0 Update 3*** or newer
2. One virtual NIC and vSwitch per NG Firewall Interface

Download the NG Firewall Virtual Machine

• Download the NG Firewall Virtual Appliance:

  1. Log into your Edge Threat Management account.
  2. Click GET STARTED at the top right-hand corner.
  3. Select the latest version and download the ISO file.

Deploy Image to ESX Server

• Once the image is downloaded, open your VMware vSphere Client and login to your server. vCenter Login.
• Once you are logged in, click **File > Deploy OVF Template…** vCenter File->Deploy.
• In the **Deploy OVF Template** wizard mark **Deploy from file:** And hit **Browse….** vCenter Deploy Wizard 1.
• Browse to the location where you saved your image and click **Open**.
• Then hit **Next**. vCenter Deploy Wizard 2.
• Read The Template Details and click **Next**. vCenter Deploy Wizard 3.
• In the "**Name and Location screen**" you may either change the name or leave it at the default. Click **Next**. vCenter Deploy Wizard 4.
• In the "**Resource Pool screen**" If you use Resource Pools, select the appropriate pool for the new NG Firewall VM and click **Next**.

  > **Note:** You can always move the VM to another Resource Pool after it's installed. vCenter Deploy Wizard 5.

- In the "**Datastore screen**" Select what datastore you want use click **Next**. vCenter Deploy Wizard 6.
- In the "**Ready to Complete**" screen, verify that everything looks OK and click **Finish**. vCenter Deploy Wizard 7.
- Wait for the **Deploying** Progress Meter. vCenter Deploy Progress Meter.
- When it is done, Click **Close**. vCenter Deployment Completed.

Verify/Configure Physical NIC to vSwitch mappings

- Setup/confirm your vSwitch Settings. Click on the **ESX host**, then select **Configuration** tab and **Hardware** > **Networking**.
- It is best practice to place your **Management Network** is on a own vSwitch. (This is not a Must but if you can make sure that NG Firewall does not exist on the same vSwitch as any Management Interface)
- On the vSwitches that NG Firewall will connect to activate **promiscuous mode** click on **Properties…**. vCenter vSwitch Properties.
- Ensure that Promiscuous has status **Accept** otherwise hit **Edit** and go to the **Security** Tab and change **Reject** to **Accept**. You will need to do this on all vSwitches that NG Firewall Virtual Machine connects to vCenter vSwitch Properties2.

Configure the Virtual Machine for your Network

- Right click on the new **Virtual Machine** and select vCenter Edit Settings.
- You will need to add new virtual NICs and connect them to the appropriate vSwitches.

  ⚠️ **Warning:** Two Bridged Interfaces to the same vSwitch will crash your ESX server. Each NG Firewall NIC should be connected to its own vSwitch. Each vSwitch should be connected to it's own Physical NIC, or at least be separated by VLAN tagging at the physical NIC level.

- In this example, you can see that the new NICs are connected to different vSwitches labeled LAN and DMZ. vCenter VM properties.
- Under **Options > VMware Tools** make sure to check the **Synchronize guest time with host** and click **OK** (vCenter VM properties/tools options).

Celebrate! You're at the End

Now you are ready to Power on your NG Firewall VM.

More Information and Troubleshooting

For more information on the underlying issues, see the following:

- Kernel documentation
- VMware documentation
- Microsoft Virtual Server documentation
- Edge Threat Management Community Support
- Edge Threat Management Live Support

For information about using your new NG Firewall software, see our NG Firewall User Guide.

1.6 Recovery Utilities

Recovery Utilities provides several administrative operations for Arista. It can only be launched while physically at the Arista box using the button or the special boot options.

- **Remap Interfaces**: This options remaps the mapping between "Interfaces" and "NICs" so that the physical network cards can be mapped to their desired use/configuration.

- **Configure Interface**: This option allows for basic configuring of an interface. Once the appliance is online and reachable further configuration is possible through the web-based administration interface.
- **Ping**: A basic ping test for testing network connectivity.
- **Upgrade**: Launch an upgrade process that will upgrade all software if upgrades are available.
- **Reboot**: Reboot the server.
- **Shutdown**: Shutdown the server.
- **Reset to Factory Defaults**: Delete all settings and return to the factory defaults.

Web Administration

Contents

2.1   Administration Interface

The Administration Interface is the main interface used to configure NG Firewall.

A registration and welcome message is displayed upon the first visit to the administration interface. It will make a suggestion about the suggested applications that may be useful for your network. You can choose to install the recommended apps or to install the apps manually.

There are four main tabs in the administration interface in the main menu:

1. **Dashboard**
2. **Apps**
3. **Config**
4. **Reports** (only visible if the Reports app is installed.)

In the sub-menu there are four views:

1. **Sessions**
2. **Hosts**
3. **Devices**
4. **Users**

Tip: Using Mozilla Firefox or Google Chrome browsers is recommended for administration.

2.2 Dashboard

The *Dashboard* provides an overview of the state of your NG Firewall. It is extremely useful for quickly viewing or monitoring what is happening on the network and the current status of the NG Firewall server.
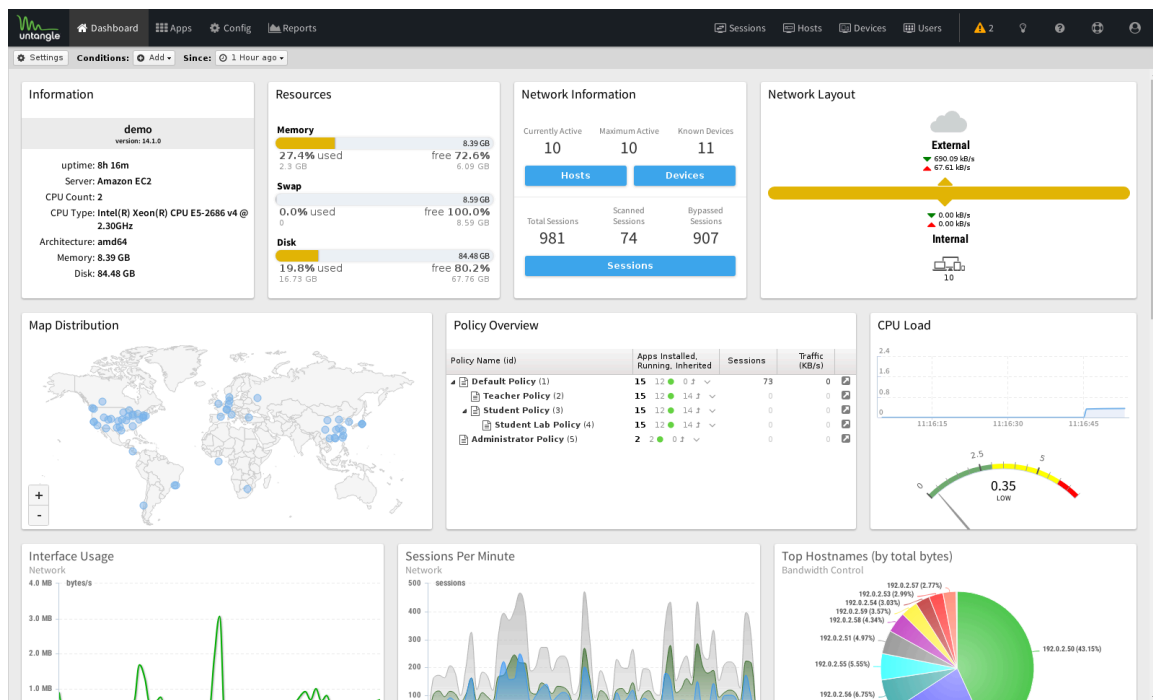
Figure 19: NG Firewall Dashboard

By default the dashboard will show several *widgets* with varying pieces of information. However, the dashboard is completely customizable. Widgets can be removed and added so the administrator sees exactly the information that is important to them on the dashboard.

There are many different type of *widgets* available:

Name	Information
Information	Shows some information about NG Firewall, like name, model, version, etc.
Resources	Shows an overview current memory swap and disk usage.
CPU Load	Shows a graph of recent CPU load.

Name	Information
Network Information	Shows an overview of the network information like session count and device/host count.
Network Layout	Shows an overview of the network layout based on the interface configuration.
Map Distribution	Shows the current sessions mapped geolocation on a world map, sized by throughput.
Report	Shows any Report Entry from Reports

To change what is displayed on the dashboard click on **Manage Widgets** at the top. From here you can show or hide the built-in widgets, or add new widgets from Reports by clicking on the **Add** button.

When adding a Report widget, you must also specify a timeframe (the number of hours worth of data to display) and a refresh interval (how often the widget refreshes on the dashboard).

When viewing a Report Entry in Reports you can also easily add it to your dashboard by clicking on the **Add to Dashboard** button.

If you see an alert icon near the top of the dashboard, hover over it for more information - these are Administrative Alerts designed to help you keep your NG Firewall healthy.

2.3 Administration Notifications

Administration Notifications appear as an exclamation point icon at the top of the rack when logged into the Administration interface or in the "Notifications" widget on the dashboard.

Overview

When logging in, the server will runs a series of tests which can take a few minutes and then it will display the administration alert icon if there are any alerts. Tests are only performed on login, to force a retest just refresh the browser or click refresh on the Notification widget on the dashboard.
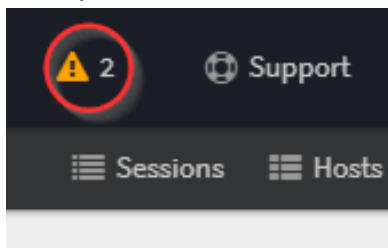
Notifications are displayed to alert the administrator of common misconfigurations or issues.

Notifications

Text	Description
Upgrades are available and ready to be installed.	The server detected software upgrades that have not been applied. Upgrades can be applied in **Config > Upgrade**.

Text	Description
DNS connectivity failed: **DNS Server IP**	The specified server's DNS settings is not providing DNS resolution. Check DNS settings of your WAN interfaces in **Config > Network > Interfaces**. It is recommended to use your ISP's DNS servers.
Failed to connect to Arista. [**address:port**]	Arista failed to successfully connect to the Arista servers. Check your network setting to make sure they are valid and that Arista is online. Also check there is no firewall between Arista and the internet that could be blocking connectivity. Arista requires an active connection to the internet for proper operation.
Free disk space is low. *[**xx%** free]*	Free disk space is running low. Contact Arista support for help determining what is using disk space and what to do about it. Please note that our recommended minimum hard disk size is at least **80** Gigs.
Disk errors reported.*Error text*	The disk (hard drive) returned some errors for certain commands. This usually means the disk has bad sectors which are non-responsive. In this case the disk (hard drive) should be immediately replaced.
Rack Name contains two or more **Application 1**	The given rack contains two or more instances of the same application. While possible this is never desired as it decreases performance and increases management complexity. Remove one of the duplicate applications.
Rack Name contains redundant apps: **Application 1** and **Application 2**	Some applications in Arista are redundant and should not both be installed in the same rack at the same time. For example, Spam Blocker is a super-set to Spam Blocker Lite. If both are run no additional spam will be blocked, but messages will be scanned twice incurring a performance hit. Remove the redundant application.
Bridge (**Interface 1** <-> **Interface 2**) may be backwards. Gateway (**Gateway IP**) is on **Interface 2**.	Often bridges can be plugged in with the WAN interfaces on the LAN and the LAN interface on the WAN. This works and passes traffic, however several applications do not behave as expected. If shown, NGFW has detected that the gateway for the main bridge interface is not on the expected interface. It is recommended to go into **Config > Network > Interfaces** and unplug each interface one at a time and verify and correct the mapping of interfaces by swapping cables around.
Interface 1 interface NIC has a high number of *RX/TX* errors.	This indicates that *ifconfig* shows a high number of RX or TX errors on the given interface card. This is typically a network layer or NIC issue. If possible, try another NIC or different duplex setting in **/admin/index.do#config/network/advanced/network_cards**.

Text	Description
Spam Blocker [Lite] is installed but an unsupported DNS server is used	Spam Blocker and Spam Blocker Lite rely on DNSBL (DNS blacklists) to categorize spam. Several publicly available and often used DNS servers do not supply access to these services. For example, google (**8.8.8.8, 8.8.4.4**), open dns (**208.67.222.222**, **208.67.222.220**), level 3 (**4.2.2.1,4.2.2.2**) do not provide resolution for DNSBL queries. It is recommended to configure Arista to use your ISPs DNS servers for effective spam filtering. If spam filtering is not required simply uninstall the spam filtering application from the rack.
Spam Blocker [Lite] is installed but a DNS server *(X, Y)* fails to resolve DNSBL queries.	This means one of the configured DNS servers does not properly resolve DNSBL queries. This will greatly degrade Spam Blocker and Spam Blocker Lite's ability to detect spam. Try configuring a different DNS server. To test this manually run `host 2.0.0.127.zen.spamhaus.org your_DNS_server` in the terminal where "your_DNS_server" is the IP of your DNS server. If it does not return results then DNSBL queries are not being properly resolved by that server.
Web Filter is installed but a DNS server (**X, Y**) fails to resolve categorization queries.	This means one of the configured DNS servers does not properly resolve Web Filter category queries. Web Filter uses DNS to query for the categorization of unknown sites. If the configured DNS servers do not properly respond to categorization queries then Web Filter will not function correctly and may slow web traffic significantly.
A DNS server responds slowly. (**X,Y,Z**) This may negatively effect Web Filter performance.	This means the specified DNS server (**Y**) on interface (**X**) responded slowly (in **Z** milliseconds) to a Web Filter categorization request. Web Filter will automatically request categorization of unknown and never before seen URLs. If DNS is performing poorly Web Filter categorization will also be slow and may negatively effect web traffic latency as Web Filter categorizes websites.
Event processing is slow (**x** ms).	Event logging is slow. This is shown when event logging takes more than 15ms on average. This can be caused by a slow disk or an extremely busy server. If you see this message, you can try a couple things. 1. Use a faster disk/disk controller to the daemon is able to more quickly write events. 2. Create less events by turning off apps and/or bypassing traffic that need not be scanned.

Text	Description
Event processing is delayed (*x* minute delay).	The event logging daemon that logs events to the database is behind. This happens when "events" are happening quicker than the events can be written to the database. This can be caused by a slow disk or a busy network. Events will be stored in queued in memory until they can be written to the disk. If the time it takes for an event to happen to be logged to the database reaches a time greater than 10 minutes this warning appears. This is not necessarily an issue, but the administrator should be aware when viewing reports and events that they will be delayed by x minutes. You can try a few things to resolve this alert: 1. Use a faster disk/disk controller to the daemon is able to more quickly write events. 2. Create less events by turning off apps and/or bypassing traffic that need not be scanned.
Packet processing recently overloaded	This warning means that at "*nf_queue: full at * entries, dropping packets(s)*" was found in `/var/log/kern.log.`" This means packets were incoming faster than the server was able to handle them. This usually indicates some misconfiguration or performance issue, or that some traffic needs to be bypassed. This can also indicate that the server is undersized for the current task and is short on memory (swapping) or disk I/O throughput or processing power. For further help with this alert, contact Arista support.
The shield is disabled. This can cause performance and stability problems.	The shield is disabled in **Config > System > Shield** > > . While sometimes useful for testing, this configuration will cause performance and stability problems. To fix verify that **Enable Shield** is checked.
Route to unreachable address: *1.2.3.4*	A static route exists in **Config** > **Network** > **Routes** > > , but the next hop is unreachable. All traffic for this route will be dropped.
Currently the number of devices significantly exceeds the number of licensed devices. (*x* > *y*)	The number of devices for which NGFW has recently processed traffic (*x*) is greater than the number of allowed devices (*y*) for the license existing on the NGFW server. In order to return to compliance it may be necessary to bypass devices or get a larger license. Please contact http://support@Arista.com for help.
DNS and DHCP services are not functioning.	This means that the DNS and DHCP service is not properly functioning. This is a serious issue that must be resolved in order for Arista to function properly. The usual cause of this is invalid options or syntax in **Config > Network > Advanced > DHCP & DNS**, or in the interface settings in **Config > Interfaces > Edit > DHCP Configuration > DHCP Options**.
The timezone has been changed since boot. A reboot is required.	The timezone has been reconfigured since boot-up and reboot is required at the earliest convenience.

Text	Description
An upgrade process has been interrupted.	An upgrade has been interrupted. This is usually the result of rebooting during an upgrade or using an alternate upgrade process or running multiple upgrades at once or some other similar scenario. Contact Arista support. (Please be sure Support access is enabled in **Config > System > Support**!)

2.3.1   Overview

Administration Notifications appear as an exclamation point icon at the top of the rack when logged into the Administration interface or in the "Notifications" widget on the dashboard. When logging in, the server will runs a series of tests which can take a few minutes and then it will display the administration alert icon if there are any alerts. Tests are only performed on login, to force a retest just refresh the browser or click refresh on the Notification widget on the dashboard.
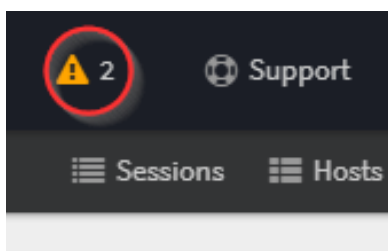
Notifications are displayed to alert the administrator of common misconfigurations or issues.

2.3.2   Notifications

Text	Description
Upgrades are available and ready to be installed.	The server detected software upgrades that have not been applied. Upgrades can be applied in Config > Upgrade.
DNS connectivity failed: **DNS Server IP**	The specified server's DNS settings is not providing DNS resolution. Check DNS settings of your WAN interfaces in Config > Network > Interfaces. It is recommended to use your ISP's DNS servers.
Failed to connect to Arista. [**address:port**]	Arista failed to successfully connect to the Arista servers. Check your network setting to make sure they are valid and that Arista is online. Also check there is no firewall between Arista and the internet that could be blocking connectivity. Arista requires an active connection to the internet for proper operation.
Free disk space is low. [**xx%** free]	Free disk space is running low. Contact Arista support for help determining what is using disk space and what to do about it. Please note that our recommended minimum hard disk size is at least **80** Gigs.
Disk errors reported.*Error text*	The disk (hard drive) returned some errors for certain commands. This usually means the disk has bad sectors which are non-responsive. In this case the disk (hard drive) should be immediately replaced.
Rack Name contains two or more **Application 1**	The given rack contains two or more instances of the same application. While possible this is never desired as it decreases performance and increases management complexity. Remove one of the duplicate applications.

Text	Description
Rack Name contains redundant apps: **Application 1** and **Application 2**	Some applications in Arista are redundant and should not both be installed in the same rack at the same time. For example, Spam Blocker is a super-set to Spam Blocker Lite. If both are run no additional spam will be blocked, but messages will be scanned twice incurring a performance hit. Remove the redundant application.
Bridge (**Interface 1** <-> **Interface 2**) may be backwards. Gateway (**Gateway IP**) is on **Interface 2**.	Often bridges can be plugged in with the WAN interfaces on the LAN and the LAN interface on the WAN. This works and passes traffic, however several applications do not behave as expected. If shown, NGFW has detected that the gateway for the main bridge interface is not on the expected interface. It is recommended to go into Config > Network > Interfaces and unplug each interface one at a time and verify and correct the mapping of interfaces by swapping cables around.
Interface 1 interface NIC has a high number of *RX/TX* errors.	This indicates that *ifconfig* shows a high number of RX or TX errors on the given interface card. This is typically a network layer or NIC issue. If possible, try another NIC or different duplex setting in `/admin/index.do#config/network/advanced/network_cards`.
Spam Blocker [Lite] is installed but an unsupported DNS server is used	Spam Blocker and Spam Blocker Lite rely on DNSBL (DNS blacklists) to categorize spam. Several publicly available and often used DNS servers do not supply access to these services. For example, google (**8.8.8.8, 8.8.4.4**), open dns (**208.67.222.222**, **208.67.222.220**), level 3 (**4.2.2.1,4.2.2.2**) do not provide resolution for DNSBL queries. It is recommended to configure Arista to use your ISPs DNS servers for effective spam filtering. If spam filtering is not required simply uninstall the spam filtering application from the rack.
Spam Blocker [Lite] is installed but a DNS server *(X, Y)* fails to resolve DNSBL queries.	This means one of the configured DNS servers does not properly resolve DNSBL queries. This will greatly degrade Spam Blocker and Spam Blocker Lite's ability to detect spam. Try configuring a different DNS server. To test this manually run `host 2.0.0.127.zen.spamhaus.org your_DNS_server` in the terminal where "your_DNS_server" is the IP of your DNS server. If it does not return results then DNSBL queries are not being properly resolved by that server.
Web Filter is installed but a DNS server (**X, Y**) fails to resolve categorization queries.	This means one of the configured DNS servers does not properly resolve Web Filter category queries. Web Filter uses DNS to query for the categorization of unknown sites. If the configured DNS servers do not properly respond to categorization queries then Web Filter will not function correctly and may slow web traffic significantly.
A DNS server responds slowly. (**X,Y,Z**) This may negatively effect Web Filter performance.	This means the specified DNS server (**Y**) on interface (**X**) responded slowly (in **Z** milliseconds) to a Web Filter categorization request. Web Filter will automatically request categorization of unknown and never before seen URLs. If DNS is performing poorly Web Filter categorization will also be slow and may negatively effect web traffic latency as Web Filter categorizes websites.

Text	Description
Event processing is slow (*x* ms).	Event logging is slow. This is shown when event logging takes more than 15ms on average. This can be caused by a slow disk or an extremely busy server. If you see this message, you can try a couple things. 1. Use a faster disk/disk controller to the daemon is able to more quickly write events. 2. Create less events by turning off apps and/or bypassing traffic that need not be scanned.
Event processing is delayed (*x* minute delay).	The event logging daemon that logs events to the database is behind. This happens when "events" are happening quicker than the events can be written to the database. This can be caused by a slow disk or a busy network. Events will be stored in queued in memory until they can be written to the disk. If the time it takes for an event to happen to be logged to the database reaches a time greater than 10 minutes this warning appears. This is not necessarily an issue, but the administrator should be aware when viewing reports and events that they will be delayed by x minutes. You can try a few things to resolve this alert: 1. Use a faster disk/disk controller to the daemon is able to more quickly write events. 2. Create less events by turning off apps and/or bypassing traffic that need not be scanned.
Packet processing recently overloaded	This warning means that at "*nf_queue: full at * entries, dropping packets(s)*" was found in `/var/log/kern.log.`" This means packets were incoming faster than the server was able to handle them. This usually indicates some misconfiguration or performance issue, or that some traffic needs to be bypassed. This can also indicate that the server is undersized for the current task and is short on memory (swapping) or disk I/O throughput or processing power. For further help with this alert, contact Arista support.
The shield is disabled. This can cause performance and stability problems.	The shield is disabled in Config > System > Shield. While sometimes useful for testing, this configuration will cause performance and stability problems. To fix verify that **Enable Shield** is checked.
Route to unreachable address: *1.2.3.4*	A static route exists in Config > Network > Routes, but the next hop is unreachable. All traffic for this route will be dropped.
Currently the number of devices significantly exceeds the number of licensed devices. (*x* > *y*)	The number of devices for which NGFW has recently processed traffic (*x*) is greater than the number of allowed devices (*y*) for the license existing on the NGFW server. In order to return to compliance it may be necessary to bypass devices or get a larger license. Please contact http://support@Arista.com for help.
DNS and DHCP services are not functioning.	This means that the DNS and DHCP service is not properly functioning. This is a serious issue that must be resolved in order for Arista to function properly. The usual cause of this is invalid options or syntax in Config > Network > **Advanced** > DHCP & DNS, or in the interface settings in Config > Interfaces > **Edit** > **DHCP Configuration** > **DHCP Options**.
The timezone has been changed since boot. A reboot is required.	The timezone has been reconfigured since boot-up and reboot is required at the earliest convenience.

Text	Description
An upgrade process has been interrupted.	An upgrade has been interrupted. This is usually the result of rebooting during an upgrade or using an alternate upgrade process or running multiple upgrades at once or some other similar scenario. Contact Arista support. (Please be sure Support access is enabled in **Config > System > Support**!)

2.4 Event Definitions

All event data is stored in the Mail messages in a relational database. As Arista and applications process traffic they create Event objects that add and modify content in the database. Each event has it's own class/object with certain fields that modify the database in a certain way.

The list below shows the classes used in the event logging and the attributes of each event object. These can be used to add alerts in Reports or for other event handling within Arista.

SpamLogEvent

These events are created by Spam Blocker and update the Database Schema table when an email is scanned.

Attribute Name	Type	Description getAction
action	SpamMessageAction	The action *getClass*
class	Class	The class name *getClientAddr*
clientAddr	InetAddress	The client address *getClientPort*
clientPort	int	The client port getMessageId
messageId	Long	The message ID *getPartitionTablePostfix* *getReceiver*
receiver	String	The receiver *getScore*
score	float	The score *getSender*
sender	String	The sender *getServerAddr*
serverAddr	InetAddress	The server address *getServerPort*
serverPort	int	The server port *getSmtpMessageEvent*
smtpMessageEvent	SmtpMessageEvent	The parent SMTP message event isSpam
isSpam	boolean	True if spam, false otherwise getSubject
subject	String	The subject getTag getTestsString
testsString	String	The tests string from the spam engine *getTimeStamp*

Attribute Name	Type	Description getAction
timeStamp	Timestamp	The timestamp *getVendorName*
vendorName	String	The application name

SpamSmtpTarpitEvent

These events are created by Spam Blocker and inserted to the Database Schema table when a session is tarpitted.

Attribute Name	Type	Description getIPAddr
IPAddr	InetAddress	The IP address *getClass*
class	Class	The class name *getHostname*
hostname	String	The host name *getPartitionTablePostfix getSessionEvent*
sessionEvent	SessionEvent	The session event *getSessionId*
sessionId	Long	The session ID *getTag getTimeStamp*
timeStamp	Timestamp	The time stamp *getVendorName*
vendorName	String	The application name

PrioritizeEvent

These events are created by the Bandwidth ControlDatabase Schema and update the table when a session is prioritized.

Attribute Name	Type	Description getClass
class	Class	The class name *getPartitionTablePostfix getPriority*
priority	int	The priority *getRuleId*
ruleId	int	The rule ID *getSessionEvent*
sessionEvent	SessionEvent	The session event *getTag getTimeStamp*
timeStamp	Timestamp	The timestamp

VirusFtpEvent

These events are created by Virus Blocker and update the Database Schema table when Virus Blocker scans an FTP transfer.

Attribute Name	Type	Description getAppName
appName	String	The name of the application *getClass*

Attribute Name	Type	Description getAppName
class	Class	The class name **getClean**
clean	boolean	True if clean, false otherwise **getPartitionTablePostfix getSessionEvent**
sessionEvent	SessionEvent	The session event **getTag getTimeStamp**
timeStamp	Timestamp	The timestamp **getUri**
uri	String	The URI **getVirusName**
virusName	String	The virus name, if not clean

VirusHttpEvent

These events are created by Virus Blocker and update the Database Schema table when Virus Blocker scans an HTTP transfer.

Attribute Name	Type	Description getAppName
appName	String	The name of the application **getClass**
class	Class	The class name **getClean**
clean	boolean	True if clean, false otherwise **getPartitionTablePostfix getRequestLine**
requestLine	RequestLine	The request line **getSessionEvent**
sessionEvent	SessionEvent	The session event **getTag getTimeStamp**
timeStamp	Timestamp	The timestamp **getVirusName**
virusName	String	The virus name, if not clean

VirusSmtpEvent

These events are created by Virus Blocker and update the Database Schema table when Virus Blocker scans an email.

Attribute Name	Type	Description getAction
action	String	The action **getAppName**
appName	String	The name of the application **getClass**
class	Class	The class name **getClean**
clean	boolean	True if clean, false otherwise **getMessageId**
messageId	Long	The message ID **getPartitionTablePostfix getTag getTimeStamp**

Attribute Name	Type	Description getAction
timeStamp	Timestamp	The timestamp *getVirusName*
virusName	String	The virus name, if not clean

FirewallEvent

These events are created by Firewall and update the Database Schema table when a firewall rule matches a session.

Attribute Name	Type	Description getBlocked
blocked	boolean	True if blocked, false otherwise *getClass*
class	Class	The class name *getFlagged*
flagged	boolean	True if flagged, false otherwise *getPartitionTablePostfix getRuleId*
ruleId	long	The rule ID *getSessionId*
sessionId	Long	The session *ID getTag getTimeStamp*
timeStamp	Timestamp	The timestamp

2.4.1   Events

Events control the handling of "events" in NG Firewall.

Within NG Firewall and the apps, when noteworthy actions occur an "event" is logged. An event is an object that describes an action. For example, an HttpRequestEvent is logged when a client on the network makes an HTTP Request. A SessionEvent is logged when a PC creates a network connection.

The Event Definitions page details all of the events and the attributes.

The platform and all apps log events through the Event Manager. The Event Manager will do several things with each event:

1. Evaluate Alert Rules below section and create, log, and send an alert if necessary.
2. Evaluate Trigger Rules from the below section and take an action if necessary.
3. Evaluate Syslog Rules from the below section and send a syslog message if necessary.
4. Send the event to Reports, if installed, for saving in the reports database.

Alerts

Alert rules are evaluated on all events logged and will log and/or alert the administrator when interesting or noteworthy events occur.

Unlike most rules, all Alert rules are evaluated even beyond the first matching rule.

Each logged event is represented by a JSON object. As each event is logged to the database the alert rules are evaluated. If all of an alert rule's conditions match the logged event the action(s) configured in the alert rule is performed.

- **Enable** determine if the alert rule is enabled.
- **Class** is the type of event this rule matches. Selecting the *Class* will determine what *Fields* are available in the conditions.

- **Conditions** lists the fields within the event object to be checked. If all of the conditions match then the rule will match.
- **Enable Thresholds** limits the alert from firing until it reaches a certain frequency threshold.
  - **Exceeds Threshold Limit** is the frequency limit for which this condition will match. If the frequency is greater than this value, then the threshold conditions matches.
  - **Over Timeframe** defines the time range, in seconds, to use to compute the frequency.
  - **Grouping Field** defines how to group thresholds by an attribute field in the events. This field is optional.
- **Log Alert** logs the event to the *Alert Event Log*
- **Send Alert** sends an email to all administrators' emails describing the event.
  - **Limit Send Frequency** limits the number of times a rule can send an alert email *To once per* the configured amount of minutes. For some cases, like a low disk space alert, this is useful to limit the number of alerts sent so that an alert is not sent every minute.

If *Exceeds Threshold Limit* is **100** and *Over Timeframe* is **60**, then the threshold condition will only match when this rules other conditions match approximately 100 times over any 60 second period. If *Group Field* is set to "CClientAddr", then the threshold load is grouped by "CClientAddr" value in the event objects. Using the above example, this would mean that the alert would only fire when a specific "CClientAddr" like "***192.168.1.100***" does something over 100 times within 60 seconds. The threshold value for other clients like "***192.168.1.150***" is tracked separately.

Adding Alert Rules

Writing and designing alert rules can be a bit of an art.

Start by finding an event that describes the action that you want to be alerted about. The Event Definitions describes all the event objects and the associated attributes on each object.

Set the *Class* to the event you want to alert about, then add conditions that check the fields to look for the events you are interested in.

As an example, lets say we want to setup an alert for when a specific website is visited by a specific user.

As a *Class* select *HttpRequestEvent* Then as a field add *domain = example.com* and *sessionEvent.username = example_user*.

We want to know if this user visits this website a single time, so we do not want to add a threshold. We want it to log this alert so we want to check *Log* and we want to send an email so we're going to check *Send Email*.

However, in reality when a user visits a website it does a lot of separate HTTP requests to load all of the components. We do not want to receive 20 emails each time a user visits a single page on that website. As we want to check *Limit Send Frequency* to *20* minutes so that we aren't flooded with emails.

There are many other alert rules that are not enabled by default that can provide some common examples.

Triggers

Triggers are similar to Alert rules, however instead of alerting when something interesting happens, trigger rules can "tag" a specific host, device, or user for a specific period of time.

Unlike most rules, all Trigger rules are evaluated even beyond the first matching rule.

This allows the system to keep state on the different hosts on the network, which can serve several purposes. For example, you can tag a specific host/device/user as using a specific application when the application is used.

There are several rules that are included but not enabled to provide some examples.

- **Enable** determine if the alert rule is enabled.
- **Class** is the type of event this rule matches. Selecting the *Class* will determine what *Fields* are available in the conditions.
- **Conditions** lists the fields within the event object to be checked. If all of the conditions match then the rule will match.
- **Enable Thresholds** limits the alert from firing until it reaches a certain frequency threshold.

  - **Exceeds Threshold Limit** is the frequency limit for which this condition will match. If the frequency is greater than this value, then the threshold conditions matches.
  - **Over Timeframe** defines the time range, in seconds, to use to compute the frequency.
  - **Grouping Field** defines how to group thresholds by an attribute field in the events. This field is optional.

- **Action Type** determines the action taken

  - **Tag Host** will tag the specified host with the specified tag.
  - **Untag Host** will remove the specified tag from the specified host.
  - **Tag User** will tag the specified user with the specified tag.
  - **Untag User** will remove the specified tag from the specified user.
  - **Tag Device** will tag the specified device with the specified tag.
  - **Untag Device** will remove the specified tag from the specified device.
- - **Target** identifies the specific host/device/user. If its a single attribute name for example 'cClientAddr' it will look up to 3 layers deep within an object for any attribute named cClientAddr. If it is a fully qualified name like 'sessionEvent.cClientAddr' it will look at that specific attribute within the specified sub-object.
  - **Tag Name** specified the string (name) of the tag to be given or removed.
  - **Tag Lifetime** specifies the lifetime of the tag when adding a tag. After the lifetime expires the tag will disappear.

Syslog

Syslog sends events via syslog messages to a remote syslog server. To use syslog simply install a syslog *receiver* on another server, then enable syslog and configure as necessary.. Some syslog products are easier to set up than others. Kiwi, a third-party syslog daemon, is a favorite of many users using Windows, while those on *nix can use rsyslog.

- **Host**: The host name or IP address of the Syslog daemon that is authorized to receive syslog messages from the NG Firewall server. Do **not** set the Host to the NG Firewall itself - this will result in the hard drive filling up very quickly and most likely crashing the box.
- **Port**: The UDP port to send syslog messages to the syslog daemon. 514 is the default as this is the default syslog port.
- **Protocol**: The protocol to use to send syslog messages. The default is UDP.

Syslog Rules

WARNING: Syslog can be a very expensive operation. If configured to send all (or most) events it can seriously negatively impact the performance of the server.

Syslog Rules determine which events are sent via syslog.

Unlike most rules, all Syslog rules are evaluated even beyond the first matching rule.

- **Enable** determine if the alert rule is enabled.
- **Class** is the type of event this rule matches. Selecting the *Class* will determine what *Fields* are available in the conditions.

- **Conditions** lists the fields within the event object to be checked. If all of the conditions match then the rule will match.
- **Enable Thresholds** limits the alert from firing until it reaches a certain frequency threshold.
-
    - **Exceeds Threshold Limit** is the frequency limit for which this condition will match. If the frequency is greater than this value, then the threshold conditions matches.
    - **Over Timeframe** defines the time range, in seconds, to use to compute the frequency.
    - **Grouping Field** defines how to group thresholds by an attribute field in the events. This field is optional.
- • **Remote Syslog** determines if the event is sent via syslog.

To send all events via syslog simply create one rule where *Class = All* and no conditions.

To send specific events to a syslog server configure the *Syslog Rules* to send the specific events to the syslog server.

Email Template

You can customize the content of email alerts by editing the Email Template. Items surrounded by the percent symbol represent system variables. You can use these throughout the **Subject** or **Body** of the message. The table below describes each variable.

Variable	Information
System.company	Your company name defined in Branding Manager.
Alert.description	The event description of the associated alert rule.
System.host	The Hostname of your NG Firewall system.
Event.class	The event class of the associated alert rule.
Event.summary	The event summary of the associated alert rule.
Event.values.keyvalue	The extended event details of the associated alert rule.

The preview window shows in real time how your changes to the **Subject** or **Body** will appear in the email message content.

2.5 Reports

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
Admin Logins	The number of total, successful, and failed admin logins over time.
Settings Changes	The number of settings changes over time.
Admin Login Events	All local administrator logins.
All Settings Changes	All settings changes performed by an administrator.

The tables queried to render these reports:

- Admin Logins
- Settings Changes

All Settings Changes

All Settings Changes is a report that provides a detailed view of any settings changes performed by an administrator on when upgrades are applied. This is available on all systems in the **Config** > **Administration** > **Reports tab**.

The Reports tab shows the timestamp when the change was made, username and hostname that made the change, and the settings files that were changed as a result.

By clicking the *Differences* button, you can see the exact changes that were made to the files. This uses a color coded 'diff' like feature to show the differences.

Red = Line was removed

Green = Line was added

Yellow = Line was changed

Port Forward Rule Example

The following shows an example of adding a port forward for DNS to the system.

Settings Change

First, you can see that the rule was added on *8/3/15* by the user admin from IP *10.24.24.40*. The settings file that changed was network.js with the appropriate version-*YYYY-MM-DD-time.js* file name.

By clicking the Differences button, you can see all changes. For this instance, only the DNS rule was added and the changes are recorded as shown below.

Settings Change

Related Topics

Reports & Events

Manage Reports

2.6 Applications

Applications are plugins that add functionality to your NG Firewall server - just like "apps" on an iPhone or Android device.

On the Apps tab you'll see the currently installed apps. Across the top there is a dropdown menu to switch to different **Policies.** Policies can be controlled via the Policy Manager app.

Apps can be installed by clicking the **Install Apps** button at the top. It will display the apps that can currently be installed. To install an app simply click on its icon. You can install as many apps at the same time as you like. After starting the installation of the desired apps you can click the **Done** button at the top to go back to the apps view.

Once installed the applications' settings can be configured by clicking on the **Settings** button or the icon of app, depending on the skin. Applications install with the suggested configuration which is the default settings and on/enabled in most cases. An application that is off/disabled will not process any network traffic. To enable a disabled application edit the settings and click **Enable** on the first tab inside the settings.

After clicking **Settings**, you will be presented with tabs for different settings sections, as well as typical buttons marked **OK**, **Cancel** and **Apply**. **Apply** saves any changes. **OK** saves any changes and closes the window. **Cancel** closes the window without saving settings. On the left hand side there is a **Remove** button which will remove the application from the current policy. The **Help** button will open the help for the tab currently being viewed.

NG Firewall has two types of Applications:

- **Filter Applications** All the Applications *above* the **Services** pane in the interface can have one instance per policy.
- **Service Applications** All the Applications *below* the **Services** pane are global and exist in all virtual racks.

Many networks only need one *policy* which means all traffic gets processed by the same apps and same configuration, but multiple policies (sometimes called "racks") are possible for bigger networks. For more information about running multiple racks check out the Policy Manager application.

To learn more about each application use the links below.

Filter Applications

Web Filter	Web Monitor	Virus Blocker	Virus Blocker Lite	Spam Blocker	Spam Blocker Lite
Phish Blocker	Web Cache	Bandwidth Control	Application Control	Application Control Lite	SSL Inspector
Captive Portal	Firewall	Intrusion Prevention	Threat Prevention	Ad Blocker	

Service Applications

Reports	Policy Manager	Directory Connector	Web Monitor	WAN Balancer
Captive Portal	IPsec VPN	OpenVPN	WireGuard VPN	Branding Manager
Configuration Backup	Live Support			

2.7   Devices

Devices provides a view of all current "devices" or unique MAC address on the local network(s).

Each row represents a single a device (unique MAC address) that has been seen on any LAN interface.

As NG Firewall scans and processes network traffic, the platform and many of the apps will save information about devices on the network. This information is stored in the "Device Table" and the Devices view provides a view into the device table.

2.7.1   Controls

The device view by default shows all devices and some basic information about each session.

1. **Refresh** refreshes the grid with the current active sessions.
2. **Reset View** resets the view to the default view. Any changes to the default view are saved in your local browser session.
3. **Add** can be used to manually add devices to the device table.

However, As devices are discovered they are automatically added to the device table.

1. **Export** exports the current device table to a JSON file.
2. **Import** imports a JSON file into the device table

More controls can be accessed by mousing over any column head and using the drowdown menu on the column header.

1. **Sort Ascending** sorts the selected column in ascending order.
2. **Sort Descending** sorts the selected column in descending order.
3. **Columns** allows the removal or addition of columns to the current view.
4. **Filter** provides a way to filter current data on this column with the provided value.

1. **Save** saves any changes manually made by the administrator

Unlike Sessions and Hosts, the device table is saved and permanent. The administrator can edit, modify, and save the values of the attributes for each Device.

2.7.2   Columns

Property	Description
MAC Address	The MAC address of this Device
MAC Vendor	The Vendor of the MAC address of this Device if known
Interface	The interface on which this Device was last seen
Last Hostname	The last hostname of this device that was learned automatically (via DHCP, DNS, or Directory Connector)
Hostname	The manually configured hostname for this device - this will be blank unless set by the administrator
Username	The manually configured username for this device - this will be blank unless set by the administrator
HTTP User Agent	The HTTP User Agent of this device (according to a recent HTTP request)
Last Seen Time	The last time this device was seen on the network
Tags	The tags of this device

2.8   Hosts

Hosts provides a view of all current "hosts" or unique IP address on the local network(s).

Each row represents a single a host (unique IP address) that has been seen on any LAN interface.

As NG Firewall scans and processes network traffic, the platform and many of the apps will save information about a host on the network. This information is stored in the "Host Table" and the Hosts view provides a view into the host table.

Controls

The host view by default shows all hosts and some basic information about each session. To view all the information for a session click on the session and all attributes are displayed in the property grid on the right side.

1. **Refresh** refreshes the grid with the current active sessions.
2. **Auto Refresh** toggles automatic refreshing of the grid.

3. **Reset View** resets the view to the default view. Any changes to the default view are saved in your local browser session.
4. **Filter** provides the ability to quickly filter all sessions my many key attributes.

More controls can be access by mousing over any column head and using the drowdown menu on the column header.

1. **Sort Ascending** sorts the selected column in ascending order.
2. **Sort Descending** sorts the selected column in descending order.
3. **Columns** allows the removal or addition of columns to the current view.
4. **Filter** provides a way to filter current data on this column with the provided value.

Columns

Table 2: Columns

Property	Description
MAC Vendor	The Vendor of the MAC address of this Host if known
Interface	The interface on which this Host was last seen
Creation Time	The creation time of this Host entry
Last Access Time	The last time an app or the platform accessed this Host entry
Last Session Time	The last time this host attempted to create a session
Address	The IP address of this Host
MAC Address	The MAC address of this Host if it is known
MAC Vendor	The Vendor of the MAC address of this Host if known
Interface	The interface on which this Host was last seen
Creation Time	The creation time of this Host entry
Last Access Time	The last time an app or the platform accessed this Host entry
Last Session Time	The last time this host attempted to create a session
ast Completed TCP Session Time	The last time this host successfuly completed a TCP session to a WAN address
Entitled Status	False if this host is not entitled to premium functionality because the limit is exceeded True otherwise
Active	True if this host is considered "active," False otherwise
HTTP User Agent	The HTTP User Agent of this host (according to a recent HTTP request)
Captive Portal Authenticated	True if this Host is authenticated with Captive Portal (at least one)
Tags	The tags of this Host
Tags String	The tags of this Host
Hostname	The official *Hostname* of this host
Hostname Source	The source of the official *Hostname* of this host

Property	Description
Hostname (DHCP)	The hostname of this host according to DHCP (Hosts often specify their hostname when retrieving a DHCP lease)
Hostname (DNS)	The hostname of this host according to reverse DNS
Hostname (Device)	The hostname of this host's MAC address according to Devices
Hostname (Device Last Known)	The last known hostname of this host's MAC address according to Devices
Hostname (OpenVPN)	The hostname according to OpenVPN
Hostname (Reports)	The hostname according to the Name Map in Reports
Hostname (Directory Connector)	The hostname according to Directory Connector
Username	The official *Username* associated with this host
Username Source	The source of the official *Username*
Username (Directory Connector)	The username according to Directory Connector
Username (Captive Portal)	The username according to Captive Portal
Username (Device)	The username of this host's MAC address according to Devices
Username (OpenVPN)	The username according to OpenVPN
Username (IPsec VPN)	The username according to IPsec VPN
Quota Size	The size of this host's quota (in bytes)
Quota Remaining	The amount of quota remaining (in bytes)
Quota Issue Time	The original issue time of this host's quota
Quota Expiration Time	The expiration time of this host's quota
Refill Quota	**Refill Quota** action will refill this Host's quota
Drop Quota	**Drop Quota** action will remove this Host's quota

2.8.1   Controls

The host view by default shows all hosts and some basic information about each session. To view all the information for a session click on the session and all attributes are displayed in the property grid on the right side.

1. **Refresh** refreshes the grid with the current active sessions.
2. **Auto Refresh** toggles automatic refreshing of the grid.
3. **Reset View** resets the view to the default view. Any changes to the default view are saved in your local browser session.
4. **Filter** provides the ability to quickly filter all sessions my many key attributes.

More controls can be access by mousing over any column head and using the drowdown menu on the column header.

1. **Sort Ascending** sorts the selected column in ascending order.
2. **Sort Descending** sorts the selected column in descending order.
3. **Columns** allows the removal or addition of columns to the current view.

4. **Filter** provides a way to filter current data on this column with the provided value.

2.8.2   Controls

The host view by default shows all hosts and some basic information about each session. To view all the information for a session click on the session and all attributes are displayed in the property grid on the right side.

1. **Refresh** refreshes the grid with the current active sessions.
2. **Auto Refresh** toggles automatic refreshing of the grid.
3. **Reset View** resets the view to the default view. Any changes to the default view are saved in your local browser session.
4. **Filter** provides the ability to quickly filter all sessions my many key attributes.

More controls can be access by mousing over any column head and using the drowdown menu on the column header.

1. **Sort Ascending** sorts the selected column in ascending order.
2. **Sort Descending** sorts the selected column in descending order.
3. **Columns** allows the removal or addition of columns to the current view.
4. **Filter** provides a way to filter current data on this column with the provided value.

2.9   Sessions

Sessions provides a view of the current sessions (also known as connections).

Each row represents a single a network session/and its properties.

As NG Firewall and all the apps learn more about a session, many will "attach" data to the session so it is globally visible and accessible to other apps. The Sessions view provides a view into everything that is known about each session.

The Sessions view provides a real-time view into the network, and can also provide a great debugging. The controls provide a way to view the current sessions of a specific application, host, user, website, policy, etc. This can be used to view activity or to verify that traffic is being handled properly, by the proper policy, etc.

Controls

The sessions view by default shows all active session and some basic information about each session. To view all the information for a session click on the session and all attributes are displayed in the property grid on the right side.

1. **Refresh** refreshes the grid with the current active sessions.
2. **Auto Refresh** toggles automatic refreshing of the grid.
3. **Reset View** resets the view to the default view. Any changes to the default view are saved in your local browser session.
4. **Filter** provides the ability to quickly filter all sessions my many key attributes.

More controls can be accessed by mousing over any column head and using the dropdown menu on the column header.

1. **Sort Ascending** sorts the selected column in ascending order.
2. **Sort Descending** sorts the selected column in descending order.
3. **Columns** allows the removal or addition of columns to the current view.
4. **Group by this Field** will group the session data by the selected column.
5. **Filter** provides a way to filter current data on this column with the provided value.

Columns

Property	Description
Creation Time	The creation time of the session (if scanned)
Session ID	The session ID (if scanned)
Mark	The netfilter connmark
Protocol	The protocol of the session (TCP/UDP)
Bypassed	True if the session is bypassed, False if it is scanned
Policy	The policy handling the session (if scanned)
Hostname	The hostname for the client address (if known)
NATd	True if the client address of the session was rewritten (NAT), False otherwise
Port Forwarded	True if the server address of the session was rewritten (port-forward), False otherwise
Tags	The tags attached to the session (inherited from Hosts, Devices, and Users)
Tags String	The list of all tags attached to the session.
Local Address	The IP address of the "local" (non-WAN) participant or the Client Address if no local address.
Remote Address	The IP address of the "remote" (WAN) participant or the Server Address if no remote address.
Bandwidth Control Priority	The priority of the session set by Bandwidth Control.
QoS Priority	The priority set by QoS.
Pipeline	The application processing order (pipeline) of the session (if scanned).
Client Interface	The network interface of the client (source).
Client Address (Pre-NAT)	The IP address of the client (initiator) of the session.
Client Port (Pre-NAT)	The port of the client (initiator) of the session.
Client Address (Post-NAT)	The IP address of the client (initiator) of the session post-NAT.
Client Port (Post-NAT)	The port of the client (initiator) of the session post-NAT.
Client Country	The country code of the client IP address.
Client Latitude	The latitude of the client IP address.
Client Longitude	The longitude of the client IP address.
Server Interface	The network interface of the server (destination).
Server Address (Pre-NAT)	The IP address of the server (receiver) of the session pre-NAT.
Server Port (Pre-NAT)	The port of the server (receiver) of the session pre-NAT.
Server Address (Post-NAT)	The IP address of the server (receiver) of the session.
Server Port (Post-NAT)	The port of the server (receiver) of the session.
Server Country	The country code of the server IP address.

Property	Description
Server Latitude	The latitude of the server IP address.
Server Longitude	The longitude of the server IP address.
Speed (KB/s) Client	The data rate of data sent by the client (updated every 60 seconds).
Speed (KB/s) Server	The data rate of data sent by the server (updated every 60 seconds).
Speed (KB/s) Total	The data rate of session (updated every 60 seconds).
Application Control Lite Protocol	The protocol according to Application Control Lite.
Application Control Lite Category	The category according to Application Control Lite.
Application Control Lite Description	The description of the protocol according to Application Control Lite.
Application Control Lite Matched?	True if Application Control Lite matched the session.
Application Control Protochain	The protochain of Application Control
Application Control Application	The application of Application Control
Application Control Category	The category of the application of Application Control
Application Control Detail	The detail of the application of Application Control
Application Control Confidence	The confidence of the match of Application Control
Application Control Productivity	The productivity of the application of Application Control
Application Control Risk	The risk of the application of Application Control
Web Filter Category Name	The category of the last web request according to Web Filter
Web Filter Category Description	The description of the category of the last web request according to Web Filter
Web Filter Category Flagged	True if this category of the web request is flagged, False if not, null otherwise
Web Filter Category Blocked	True if this category of the web request is blocked, False if not, null otherwise
Web Filter Flagged	True if the last web request is flagged, False if not, null otherwise
HTTP Hostname	The HTTP hostname if an HTTP session.
HTTP URL	The HTTP URL of the last HTTP request of this session.
HTTP User Agent	The HTTP User Agent of the last HTTP request of this session.
HTTP URI	The HTTP URI of the last HTTP request of this session.

Property	Description
HTTP Request Method	The HTTP Request Method of the last HTTP request of this session.
HTTP Request File Name	The HTTP Request filename of the last HTTP request of this session.
HTTP Request File Extension	The HTTP Request filename extension (.exe) of the last HTTP request of this session.
HTTP Request File Path	The HTTP Request file path of the last HTTP request of this session.
HTTP Content Type	The HTTP Content Type of the last HTTP response of this session.
HTTP Referrer	The HTTP Referrer of the last HTTP request of this session.
HTTP Response File Name	The HTTP Response filename of the last HTTP response of this session.
HTTP Response File Extension	The HTTP Response filename extension (.exe) of the last HTTP response of this session.
HTTP Content Length	The HTTP content length of the last HTTP response of this session.
SSL Subject DN	The Subject DN of the SSL certificate of this session.
SSL Issuer DN	The Issuer DN of the SSL certificate of this session.
SSL Inspected	True if SSL Inspected, False if not inspected, null otherwise.
SSL SNI Hostname	The SNI hostname specified in the request of this session (if specified).
FTP Filename	The name of the last file downloaded in this session via FTP.
FTP Data Session	True if this is an FTP data session, null otherwise.

2.10   Users

Users provides a view of all current **users** or unique username on the local network(s).

The basics of users and user management are described in the User Management documentation.

Each row represents a single a user (unique username) that has been seen on any LAN interface.

As NG Firewall scans and processes network traffic, the platform and many of the apps will save information about users on the network. This information is stored in the "User Table" and the Users view provides a view into the user table.

Controls

The user view by default shows all users and some basic information about each session.

- **Refresh** refreshes the grid with the current active sessions.
- **Reset View** resets the view to the default view. Any changes to the default view are saved in your local browser session.
- **Add** can be used to manually add users to the user table.

However, As users are discovered they are automatically added to the user table.

- **Export** exports the current user table to a JSON file.
- **Import** imports a JSON file into the user table

More controls can be access by mousing over any column head and using the drowdown menu on the column header.

- **Sort Ascending** sorts the selected column in ascending order.
- **Sort Descending** sorts the selected column in descending order.
- **Columns** allows the removal or addition of columns to the current view.
- **Filter** provides a way to filter current data on this column with the provided value.
- **Save** saves any changes manually made by the administrator

Unlike **Sessions** and **Hosts**, the user table is saved and permanent. The administrator can edit, modify, and save the values of the attributes for each User.

Quotas

User bandwidth quotas are assigned through the Bandwidth Control app after configuring the Bandwidth Control setup wizard. In the Users view you can:

- View the status of user quotas.
- Override the default quota.
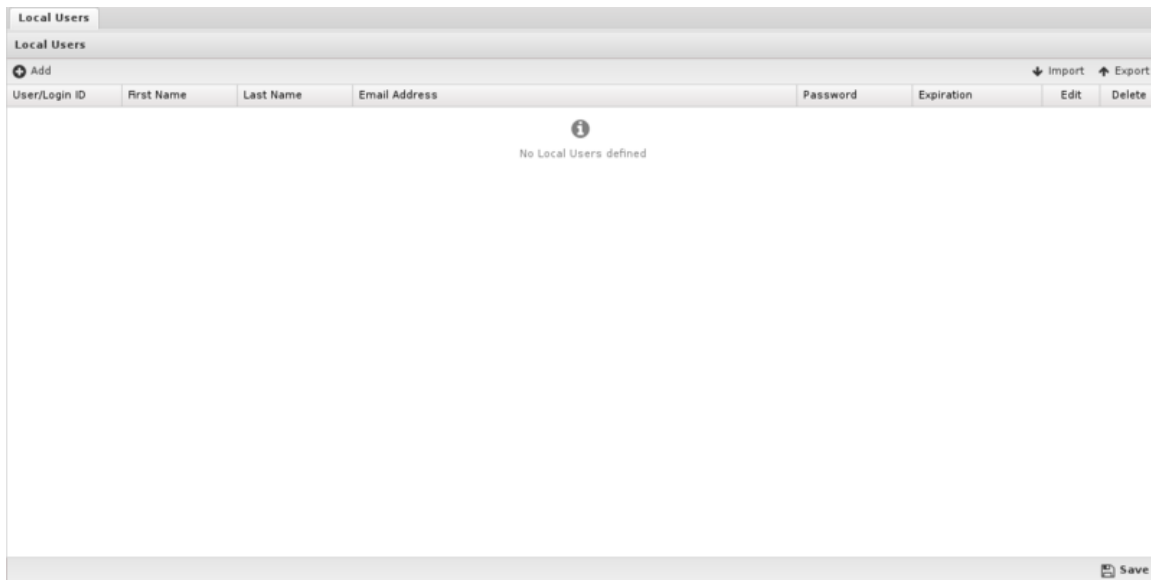- Refill the quota.
- Drop the quota.

Columns

Table 3: Columns

Property	Description
Username	The username of this entry
Creation Time	The creation time for this user entry
Last Access Time	The last time this user entry was accessed
Last Session Time	The last time this user created a network session
Quota Size	The size of this host's quota (in bytes)
Quota Remaining	The amount of quota remaining (in bytes)
Quota Issue Time	The original issue time of this host's quota
Quota Expiration Time	The expiration time of this host's quota
Quota Refill	Refill Quota action will refill this Host's quota
Quota Drop	Drop Quota action will remove this Host's quota
Tags	he tags of this user

2.11   Local Users

Local Users stores a list of users that can be used by the applications.

For example, Captive Portal and OpenVPN can use the local directory to authenticate users.

To add new users click on the Add button. You must supply a username, first name, last name, email address, and password. Only the administrator can set the password for a given user. Users can be imported or exported using the import/export buttons on the upper right.
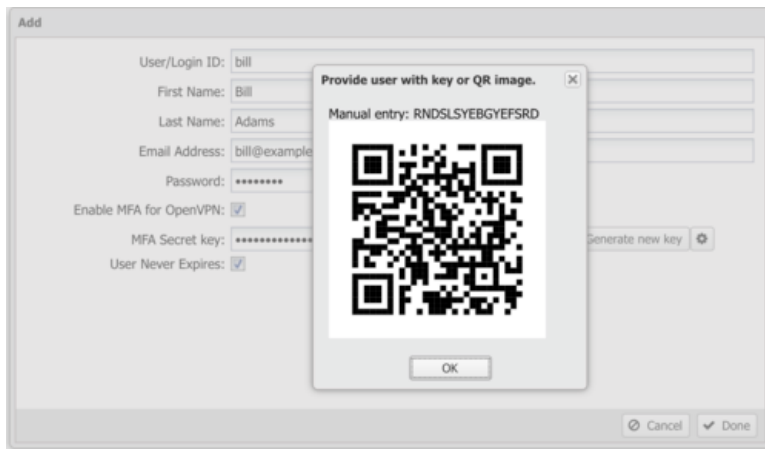
An expiration date can be specified for a user. If the expiration date has passed that user will no longer be authenticated.

To use the Local Directory, configure apps such as Captive Portal and OpenVPN to authenticate against the Local Directory while requiring user authentication.

MFA and OpenVPN

You can enable TOTP based multi factor authentication for OpenVPN client connections. When adding a user, select **Enable MFA for OpenVPN** and click **Generate new key**.

After generating a key, click the gear icon to show the QR code. Use the generated code in any TOTP mobile app such as Google Authenticator. The TOTP app generates a temporary that the user enters into their OpenVPN client. **Note:** You must also enable MFA for client configurations in OpenVPN.

> ⚠️ **Warning:** Typically, when passwords are stored, password hashes are saved and the original cleartext password is forgotten so administrators do not have access to user passwords. However, The passwords for users that are stored in the local directory are stored in cleartext because some applications and features (L2TP) depend on access to the cleartext password. Administrators do have access to cleartext user passwords and caution is advised.

2.12   Local Directory

Local Directory stores a list of users that can be used by the applications. It also supports RADIUS for 802.1x authentication from properly configured wireless network access points.

The RADIUS Server can be enabled to allow WiFi users to authenticate as any user configured in *Local Directory*.

The RADIUS Proxy can be enabled to allow WiFi users to authenticate with credentials that are validated with a configured Active Directory Server.

- Local Users
- RADIUS Server
- RADIUS Proxy
- RADIUS Log

2.13   Report Viewer

Reports provide a graphical view of the network traffic and actions of your NG Firewall. Various reports are available within applications and base system components. The reports can be manipulated to drill down, customize, and export data in many ways using the Report Viewer.

2.13.1 Reports
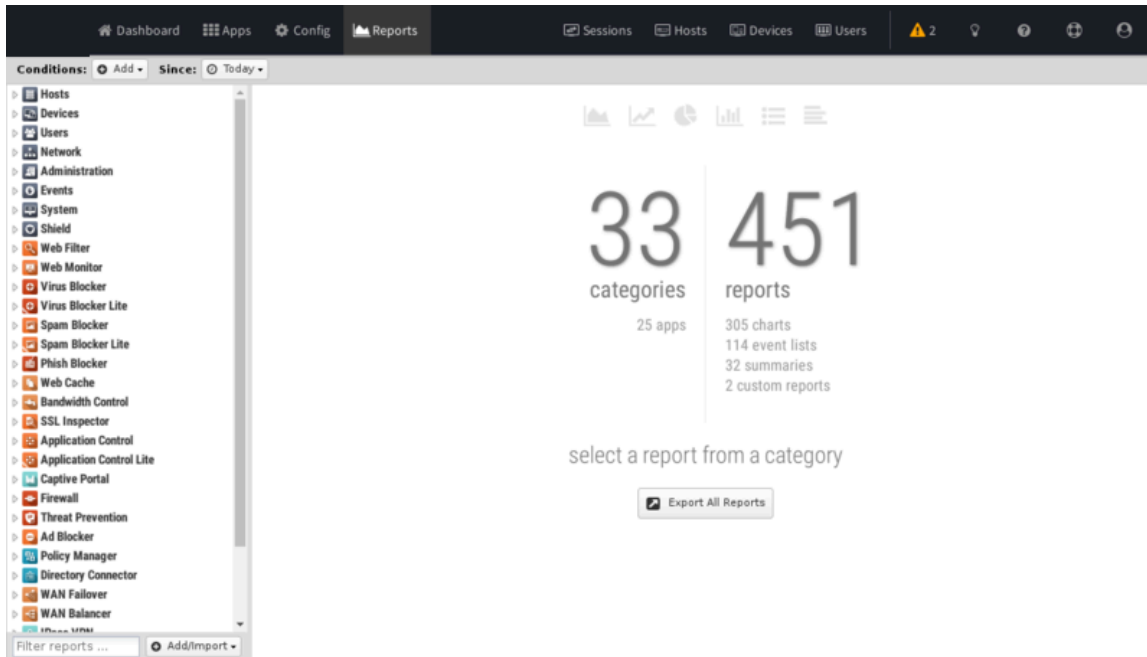
2.13.2 Report Viewer Panels

There are a few panels in the Report Viewer:

- The top panel: This top panel (just below the navigation menu) allow you specify which data is viewed. By default, there is just a timeframe and no conditions, so reports will show all data for the specified timeframe. Conditions can be viewed to view more specific data, such as a specific host, user, domain, application, web category, etc.
- The left panel: This allows you to choose the report you wish to view. At the bottom you can use the search box to quickly find reports with the specified string in the title. You can also import and create new reports using the "**Add/Import**" button.
- The chart panel: This panel shows you the specified report. It also includes several action buttons at the top.
- The data panel: The data panel, hidden by default, can be displayed by clicking on the "`Data View`" button in the chart panel. This will show the raw data used to generate the chart and allow the user to export the data by clicking the "**Export Data**" button at the bottom.

Conditions

The Conditions panel appears at the top panel and can be used to filter data displayed in reports. For example, to view a *specific* host's report, you can add a condition for Client = *192.168.1.100* and then all reports available will only show data where the client is *192.168.1.100*. Multiple conditions can be added to drill down and inspect data. Conditions can also be added quickly by clicking on slices in pie charts.

The Add Condition dropdown contains many commonly used conditions, or the full list of all tables and columns can be browsed by clicking on the *More* button to add conditions for any database column.

> **Note:** Conditions will not apply to all reports. For example, If viewing a specific users report by adding a condition where *Username* = *foobar* - many reports will be greyed out and unviewable. This is because the data used to generate those reports is not relevent to the specific user (it does not contain a username column). For example, the CPU usage report is a system report that is not relevant to a specific network user and so there is no way to filter that data by user.

Condition Operators

The second field in the condition is the logical operator that will be used in evaluating the condition value defined in the last field. In most use cases the default = operator is what you want to use. However, there are several other operators available that make the reports and alerts a whole lot more powerful.

A detailed outline of each operator is on the Operators page.

Conditions Example - Policy by Policy ID

In many cases, you may just want to see the traffic related to a specific policy within Policy Manager. This can be accomplished very easily by adding a condition using the Quick Add feature.

1. In the Conditions panel, select **Add**.
2. Choose **Policy ID** and specify equals and the **policy ID** in question.
3. The conditions is applied and will remain applied as you switch between reports.

Conditions Example - Web Filter Categories

From pie charts, you can quickly add a condition from the Current Data window. This can be handy for use with the Web Filter category selection which we'll use for this example. Once the condition is applied, we can then use other reports to drill down to find out more information about the traffic such as which user might be responsible.

1. Open Report Viewer or the Web Filter Reports tab.
2. Select the **Top Categories** report (by size or requests). In our example, you can see Games was at the top.
3. Click on the Games pie slice, and when prompted to add a condition click **Yes**.
4. All Reports can now be viewed for Games only traffic.
5. For example, the Top Clients (by request) will show the clients that visited the most gaming sites.
6. For example, the Web Usage (scanned) will show **Gaming** web usage throughout the day of the network.

2.13.3   Application Specific Report Pages

- All Reports
- Web Filter Reports

General Configuration

Contents

- Config
- About
- Administration
- Events
- Local Directory
- System
- Email

3.1   Config

The config tab holds all the settings related to configuration of the NG Firewall server itself and settings for components of the *platform* that apps may interact with.

This is an list of all sections available under the **Config** tab in the Administration UI.

Network

The *Network* configuration contains all the settings to control how your NG Firewall server routes and handles network traffic. Properly configuring network settings is critical for proper operation.

- Interfaces
- Hostname
- Services
- Port Forward Rules
- NAT Rules
- Bypass Rules
- Filter Rules
- Routes
- DNS Server
- DHCP Server
- Advanced

    - Options
    - QoS
    - Access Rules
    - UPnP
    - Network Cards
    - DNS and DHCP
    - Netflow
    - Dynamic_Routing
- Network Reports
- Troubleshooting

The Network Configuration documentation documents how networking in NG Firewall functions and is commonly configured.

Administration

Administration controls the administration-related functionality of the NG Firewall server.

- Admin
- Certificates
- SNMP
- Skins
- Google

Email

Email contains all the email-related configuration of the NG Firewall server.

- Outgoing Server
- Safe List
- Quarantine

Local Directory

Local Directory stores a list of users that can be used by the applications. It also supports RADIUS for 802.1x authentication from properly configured wireless network access points.

The RADIUS Server can be enabled to allow WiFi users to authenticate as any user configured in *Local Directory*.

The RADIUS Proxy can be enabled to allow WiFi users to authenticate with credentials that are validated with a configured Active Directory Server.

- Local Users
- RADIUS Server
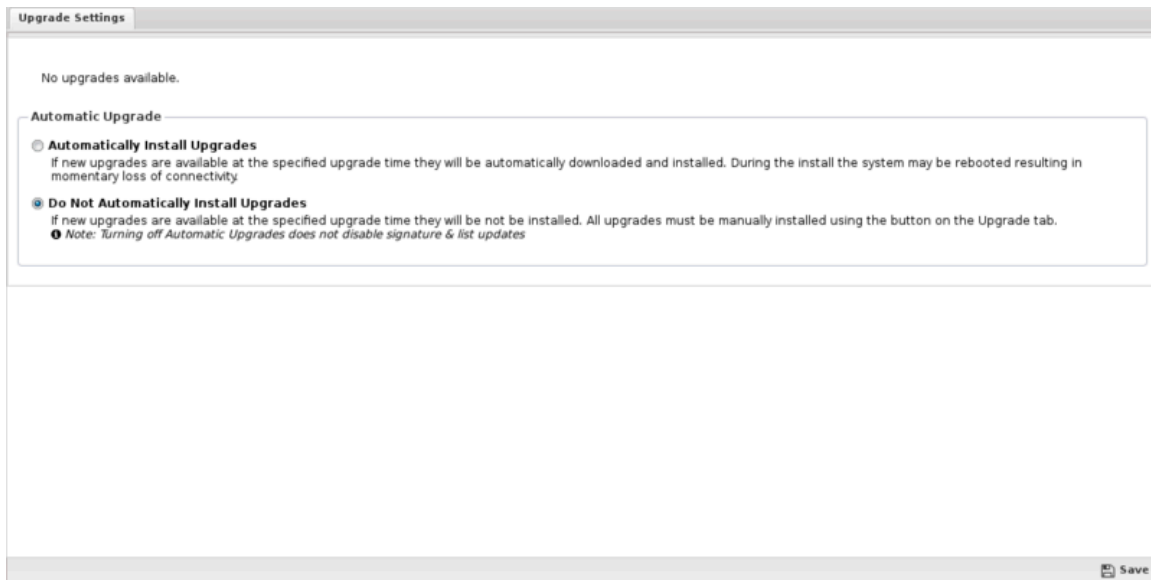- RADIUS Proxy
- RADIUS Log

Upgrade

Upgrade allows the server to upgrade and contains upgrade-related settings.

Upgrade Settings

Upgrades show the currently available upgrades if any. If upgrades are available, an upgrade can be started by pressing the *Upgrade* button at the top under Status.

To see changes see the Changelogs.

After the upgrade begins, it will download the new packages (which may take some time) and then the upgrades will be applied. Do not reboot or power off the server during the upgrade.

If *Automatically Install Upgrades* is checked, NG Firewall will automatically check for new versions and upgrade if available.

Automatic Upgrade Schedule configures when NG Firewall will automatically upgrade if upgrades are available. NG Firewall will automatically upgrade at the specified time on the days of the week than are checked.

Upgrade FAQs?

When will I get the upgrade version?

- Upgrades are rolled out over time to NG Firewall deployments. The rollout can stretch out several weeks. If you want the upgrade immediately, email your UID to the Support team and request they add your UID to the Early Upgrade list.

How do I know when the new version is available for my NG Firewall?

- In the Upgrade page of your NG Firewall, when a new version is available, the Upgrade button will appear. If automatic upgrade setting is enabled, your NG Firewall will upgrade automatically once the upgrade is available on day and time specified.

Does the upgrade require a reboot?

- If a reboot is needed, the upgrade will reboot automatically once the upgrade is installed. There is no need for a manual reboot. Most upgrades will not reboot as there is no kernel change.

How long does the upgrade take?

- It's difficult to be precise since customer platforms, Internet connection speed and complexity of the upgrade vary. Generally upgrades take less than 20 mins. If the database version is changed as part of the NG Firewall upgrade, the process will take longer as the database will need to be converted. There are extreme cases of the upgrade taking over an hour.

Do I need to reinstall?

- No, upgrade process will update all the components on the NG Firewall seamlessly.

Where can I get what is changed in the new version?

- Release changes are posted on the NG_Firewall_Changelogs page.

System

System contains settings related to the server

- Regional
- Support
- Logs
- Backup
- Restore
- Protocols
- Shield
- System Reports

About

About contains system information
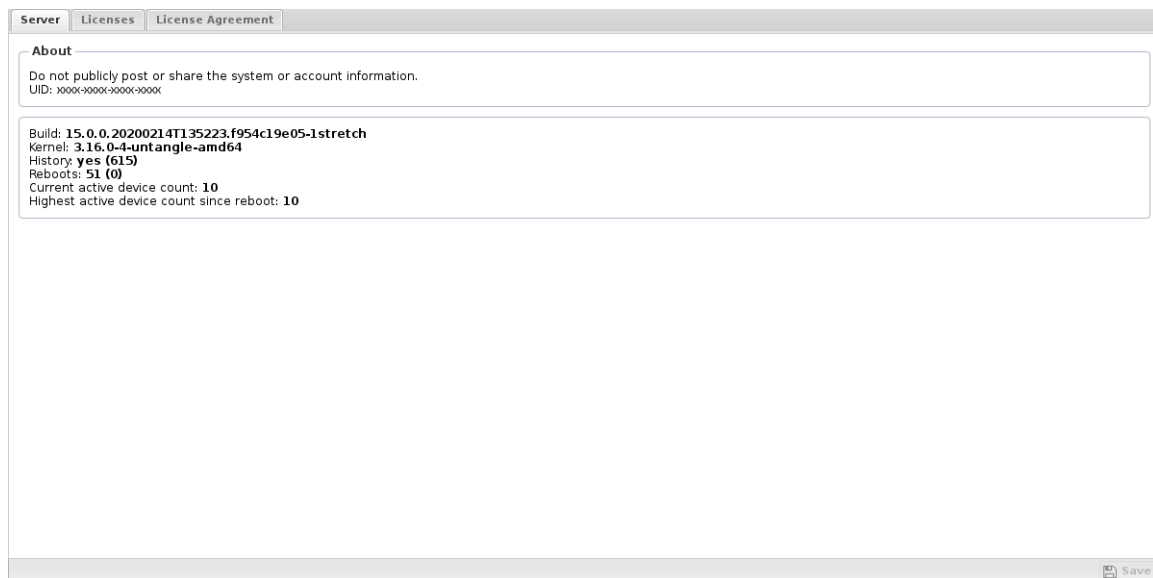
- Server
- Licenses
- License Agreement

Reports

The reports tab is only visible if the Reports app is currently installed. To read more about reports view the reports documentation.

3.2   About

Discusses Servers, Licences, and License agreements.

3.2.1   Server

This tab shows the current information about the Arista server.

Unique ID

The first field shows the Unique ID (UID) of the Arista server. The UID is a 16 alphanumeric code that uniquely identifies this server for licensing and tracking purposes.

Never share the UID of the server.

The UID is generated automatically on install and each server must have a unique UID to function properly.

Cloning servers post-installation will create two servers with identical UIDs which will result in problems and licensing issues.

Server Information

The second field shows the build version and server information.

Build shows the version of the Arista-vm.

Kernel shows the kernel version. The other fields are used by Arista support.

Current "licensed" device count shows the current number of devices in the host table that count as "licensed" devices.

Highest "licensed" device count since reboot shows the highest value of licensed devices seen by this Arista since reboot.

3.2.2   Licenses

The *Licenses* grid shows the current licenses for this Arista server (this UID). Hitting the **Refresh** button will refresh the current licenses state from the Arista licenses server.

If you have any issues with incorrect licenses or any licensing issues, contact support@arista.com so we can help resolve them.

3.2.3   License Agreement

There is currently no text in this page. You can search for this page title in other pages, search the related logs, or create this page.

3.3   Administration

Administration controls the administration-related functionality of the NG Firewall server.

- Admin

  Admin stores settings related to the administration settings for NG Firewall.
- Certificates

  NG Firewall uses digital certificates when serving web content via SSL.
- Simple Network Management Protocol

  Simple Network Management Protocol (SNMP) can be used to remotely query and monitor the current state of the ETM server.
- Skins

  Skins control the look and feel of the administration interface. Skins allow customization and tuning of the look and feel of the Arista administration.
- Reports

  Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.
- Google

Google Drive connector enables certain features of NG Firewall to store data to the connected Google account. For example, you can set the Configuration Backup app to store configuration backups to Google Drive.

3.3.1 Admin

Admin stores settings related to the administration settings for NG Firewall.

Admin Accounts

This table stores the administration accounts that can administer NG Firewall. Administrators have full administrator/root access to the NG Firewall server.

By default, there is only one admin account with the password set during the Setup Wizard. Other accounts can be created. This can be useful in a few scenarios:

- if you have multiple administrators and you wish to be able to distinguish who logged in at what time.
- you want to be able to easily disable/enable access for an administrator without changing the admin password.

Additional administrator accounts are also administrators. They also have full administrator/root access.

Allow HTTP Administration

If *Allow HTTP Administration* is checked, administration will be allowed on HTTP (unencrypted) on the primary address of non-WAN interfaces via each non-WAN interface on the port configured in Services. If unchecked then administration will not be allowed on HTTP, only on HTTPS.

> **Note:** Unchecking *Allow HTTP Administration* does not actually close the HTTP port as this service is used for other functions like blockpages, Captive Portal, etc.

> **Note:** Note: HTTPS on WANs access is controlled in **Config** > **Network** > **Advanced** > **Filter Rules** > **Access Rules**. To enable, HTTPS administration on WANs (external) check the Allow HTTPS on WANs rule.

Root

For convenience, when saving administrator account settings the "root" shell password is set to the the password of the *admin* account. If there is no *admin* account (because it was renamed to something else) then the root password is not set.

The root password is stored separately than administrator account passwords. It can be changed in the shell using passwd to any value. However if you modify the *admin* password the root password will be set to the new "admin" password when saved.

Password Recovery

If you forget the "admin" password - follow the Password Recovery process to reset the administration login/password settings to default.

3.3.2   Certificates

NG Firewall uses digital certificates when serving web content via SSL.

About Digital Certificates

The server certificate is mainly used to provide secure access to the Administrative Console, as well as the Email Quarantine features. The server also needs to generate imitation certificates on the fly when using the SSL Inspector application. There are two different ways to configure the certificate used by your server, depending on your specific requirements:

1. Create and use a server certificate signed by the internal certificate authority
2. Create a Certificate Signing Request (CSR) which you can have signed by a third party certificate authority.

If you plan on using the SSL Inspector application, **option #1** is likely a good choice. Since you'll need to install the root certificate on all client computers and devices to use SSL Inspector effectively, it makes sense to also sign the certificate used by the NG Firewall server with this same CA.

If you aren't going to use the SSL Inspector, or have some other reason to prefer a third party certificate, then **option #2** may be a better choice. This will allow you to obtain and use a server

certificate signed by a third party authority. The benefit here, assuming you use one of the standard and well known providers, is that their root certificate will already be included in the list of trusted CA's on client computers and devices, so you won't have to distribute and install a new root certificate.

Certificate Authority

During the initial server installation, a default Certificate Authority (CA) was created automatically. This CA is used to create and sign imitation certificates that are generated on the fly by the SSL Inspector application. It was also used to sign the default server certificate used by the server itself. You can use the default CA as is, or you can generate a new CA if you want to customize the information contained in the root certificate.

Generate Certificate Authority

When you click this button to generate a new CA, you will be presented with a popup form where you can enter the details to be included in the Subject DN of the new root certificate. Since this operation is creating a root certificate and not a server certificate, the CN field can contain most anything you like. Once the form is complete and you click the Generate button, the new CA will be created and the Certificate Authority information fields will be updated to display the contents of the new certificate.

Download Root Certificate

Click this button to download the root_authority.crt certificate file of the Certificate Authority on the NG Firewall server. If you are using SSL Inspector, or if you have configured your NG Firewall server to use a server certificate signed by the internal Certificate Authority, you will need to download and install this certificate on all client computers and devices to eliminate certificate warning messages when browsing or accessing secure content.

Alternatively, you can download the certificate from a client system by navigating to http://yourserver/cert.

Upload Root Certificate (CA)

This option enables you to upload the root certificate and key files you may have generated using a different source. You can either paste the contents of the certificate and key files or upload the PEM formatted files.

View other Root Certificate Authorities

This option enables you to view other Root Certificates that you may have previously uploaded. If necessary, you can revert back to a previous Root Certificate.

Server Certificate

The Server Certificate is used to secure all HTTPS connections with the NG Firewall server. This mainly applies to the Administrative Console and the Email Quarantine pages.

During the initial server installation, a default certificate is created and signed using the default Certificate Authority that was also created during installation. You can use the default root certificate as is, or you can generate a new server certificate if you want to customize the information contained in the server certificate.

Generate Server Certificate

When you click this button to generate a new server certificate, you will be presented with a popup form where you can enter the details to be included in the Subject DN of the server certificate. All fields are optional except for the Common Name (CN) field, which should contain the hostname that will be used to access the server.

Example: `hostname.domain.com`

Once the form is complete and you click the Generate button, the new server certificate will be created and the NG Firewall server will start using it immediately. The Server Certificate information fields will also be updated to display the contents of the new certificate.

Third Party Certificate

Instead of using a certificate signed by the local CA, you may instead prefer to have the NG Firewall server use a certificate signed by well-known CA such as VeriSign or Thawte. The advantage to using this type of certificate is client computers and devices will need no additional configuration, since most browsers are already configured to trust certificates issued by these authorities.

> **Note:** This has nothing to do with SSL Inspector and is just the certificate used when connecting to web services running on the NG Firewall server itself (Administration, Captive Portal, Quarantine, and so on).

Upload Server Certificate

Click the **Upload Server Certificate** button to upload an official signed certificate provided by a CA, or a certificate that you generated yourself.

Certificates from CAs are provided in many different formats. The **Import a certificate or key file** button can be used to upload the certificates and keys. First, press **Import a certificate or key file** and select the certificate. Second, press the **Import a certificate or key file** and select the private key file. Repeat this process for any additional separate intermediate certificates (not commonly required). When finished, the `Server Certificate` field should contain the server cert, and the "Certificate Key" field should contain the private key. Additionally the "Optional Intermediate Certificates" field may be populated if the CA provided an intermediate certificate. At this point click **Upload Certificate** to upload the certificate. Don't forget to adjust how the new certificate will be used (HTTPS, IPSEC, etc) in the **Server Certificates** table!

Alternatively, instead of importing files you can copy and paste the certificate, key, and intermediate certificates provided by the CA into the fields.

Create Signature Signing Request

Click the **Create Signature Signing Request** button to generate a signature signing request, you will be presented with a popup form where you can enter the details to be included in the Subject DN of the CSR. Once the form is complete and you click Generate, a server_certificate.csr file will be downloaded to your computer. The certificate authority you choose will require this file, and possibly additional information to verify that you are the "owner" of the website for which you are requesting the certificate. When they receive all the required information, and any associated fee, they will issue you a new certificate file which you can upload to the NG Firewall server.

Import Signing Request Certificate

> **Note:** Before importing your signed certificate, make sure it is formatted as a PEM file. This format is common with Apache or Linux type systems. If your certificate is not in PEM format, you need to re-download the certificate in PEM format from your certificate authority, or you can use tools such as OpenSSL to convert the certificate into the correct format.

When you receive your signed certificate, click the **Import Signing Request Certificate** button to upload the certificate to the NG Firewall server. Certificates are provided in many different formats.

You can select the **Import a certificate file** to upload a certificate file provided by the signer. This will parse the file and put the result in the displayed `Server Certificate` field and any other optional "Intermediate Certificates" in the `Optional Intermediate Certificates` field. To finish the upload click the **Upload Certificate** button.

Alternatively, you can copy paste the certificate (text) provided by the signer into the fields and click **Upload Certificate**.

### 3.3.3   Simple Network Management Protocol

Simple Network Management Protocol (SNMP) can be used to remotely query and monitor the current state of the ETM server.

If *Enable SNMP Monitoring* is checked then the SNMP daemon is enabled. Access to the SNMP daemon is controlled via the Access Rules. ETM uses snmpd and the following settings will control how the SNMP daemon is configured.

Community (Get)	This community is for a Get* operation, which is the most common method of communication. An SNMP community is the group to which devices and management stations running SNMP belong. The SNMP community defines where information is sent. The SNMP community acts as a password. ETM Server will not respond to requests from management system that do not belong to its community.
System Contact	Email address of the System Administrator that should receive SNMP messages.
System Location	Description of the system's location. Simply use the default if you don't want to specify a location.
Enable Traps	If checked, SNMP traps (events) will be sent to the configured host/port.
Community (Traps)	This community is for a Trap or Inform operation, which is a rare method of communication. An SNMP community is the group to which devices and management stations running SNMP belong. The SNMP community defines where information is sent. The SNMP community acts as a password. ETM Server will not respond to requests from management stations that do not belong to its community.
Host	The host name or IP address of the management system that is authorized to receive statistics from the ETM Server.
Port	The default port for SNMP traps is *162*.

3.3.4 Skins

Skins control the look and feel of the administration interface. Skins allow customization and tuning of the look and feel of the Arista administration.
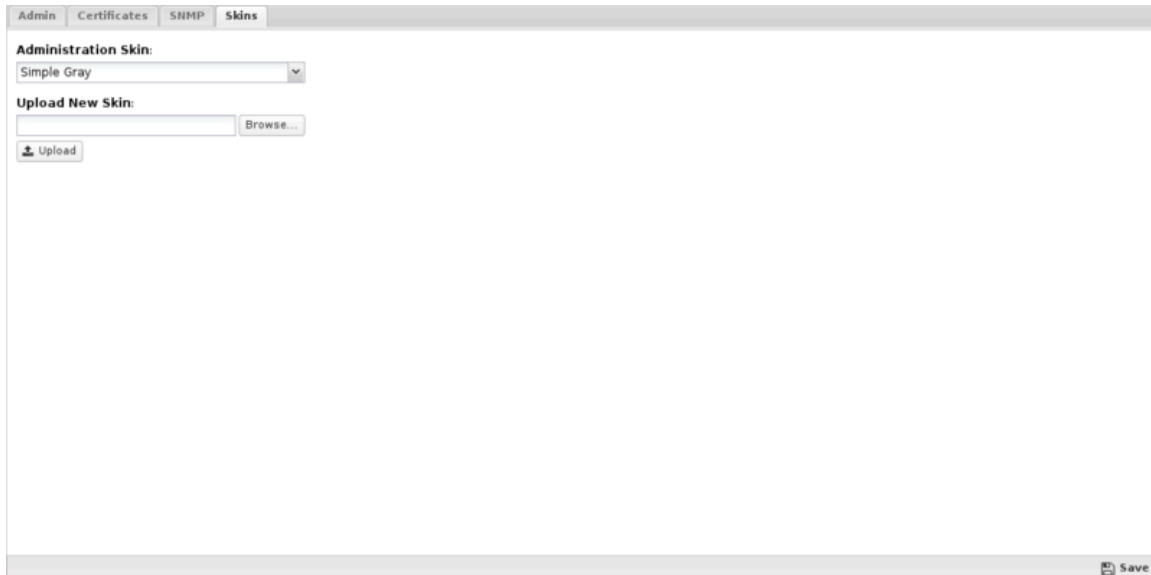
Administration Skin configures the skin to be used to render the administration UI.

Upload New Skin allows for the upload of custom skins.

> *Example:* Dell Rack AMD Skin

> Dell Rack AMD

> *Example:* Black Cammo Skin

> Black Cammo Skin

Custom Skins

Custom Skins can be created. It requires extensive work and knowledge of HTML and CSS.

Download the Arista Skins Cookbook for instructions on creating custom skins.

3.3.5 Reports

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
Admin Logins	The number of total, successful, and failed admin logins over time.
Settings Changes	The number of settings changes over time.
Admin Login Events	All local administrator logins.
All Settings Changes	All settings changes performed by an administrator.

The tables queried to render these reports:

- Admin Logins
- Settings Changes

All Settings Changes

All Settings Changes is a report that provides a detailed view of any settings changes performed by an administrator on when upgrades are applied. This is available on all systems in the **Config** > **Administration** > **Reports tab**.

The Reports tab shows the timestamp when the change was made, username and hostname that made the change, and the settings files that were changed as a result.

By clicking the *Differences* button, you can see the exact changes that were made to the files. This uses a color coded 'diff' like feature to show the differences.

Red = Line was removed

Green = Line was added

Yellow = Line was changed

Port Forward Rule Example

The following shows an example of adding a port forward for DNS to the system.

Settings Change

First, you can see that the rule was added on *8/3/15* by the user admin from IP *10.24.24.40*. The settings file that changed was network.js with the appropriate version-*YYYY-MM-DD-time.js* file name.

By clicking the Differences button, you can see all changes. For this instance, only the DNS rule was added and the changes are recorded as shown below.

Settings Change

Related Topics

Reports & Events

Manage Reports

3.3.6   Google

Google Drive connector enables certain features of NG Firewall to store data to the connected Google account. For example, you can set the Configuration Backup app to store configuration backups to Google Drive.

You can also set the Reports app to backup the database to Google Drive. Before setting these features, you must connect your NG Firewall system to a Google account.
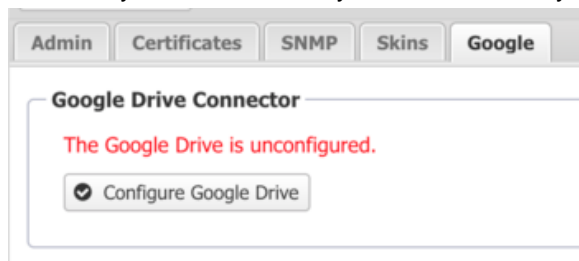
Configure Google Drive: Click this button to connect your Google account. This opens a browser window that asks you to allow your NG Firewall system to access your Google account. After you complete the setup, the status shows the connector as configured. You can *Reconfigure Google Drive* or *Disconnect Google Drive*.

3.4   Events

Events control the handling of "events" in NG Firewall.

Within NG Firewall and the apps, when noteworthy actions occur an "event" is logged. An event is an object that describes an action. For example, an HttpRequestEvent is logged when a client on the network makes an HTTP Request. A SessionEvent is logged when a PC creates a network connection.

The Event Definitions page details all of the events and the attributes.

The platform and all apps log events through the Event Manager. The Event Manager will do several things with each event:

1. Evaluate Alert Rules below section and create, log, and send an alert if necessary.
2. Evaluate Trigger Rules from the below section and take an action if necessary.
3. Evaluate Syslog Rules from the below section and send a syslog message if necessary.
4. Send the event to Reports, if installed, for saving in the reports database.

Alerts

Alert rules are evaluated on all events logged and will log and/or alert the administrator when interesting or noteworthy events occur.

Unlike most rules, all Alert rules are evaluated even beyond the first matching rule.

Each logged event is represented by a JSON object. As each event is logged to the database the alert rules are evaluated. If all of an alert rule's conditions match the logged event the action(s) configured in the alert rule is performed.

- **Enable** determine if the alert rule is enabled.
- **Class** is the type of event this rule matches. Selecting the *Class* will determine what *Fields* are available in the conditions.
- **Conditions** lists the fields within the event object to be checked. If all of the conditions match then the rule will match.
- **Enable Thresholds** limits the alert from firing until it reaches a certain frequency threshold.
    - **Exceeds Threshold Limit** is the frequency limit for which this condition will match. If the frequency is greater than this value, then the threshold conditions matches.
    - **Over Timeframe** defines the time range, in seconds, to use to compute the frequency.
    - **Grouping Field** defines how to group thresholds by an attribute field in the events. This field is optional.
- **Log Alert** logs the event to the *Alert Event Log*
- **Send Alert** sends an email to all administrators' emails describing the event.
    - **Limit Send Frequency** limits the number of times a rule can send an alert email *To once per* the configured amount of minutes. For some cases, like a low disk space alert, this is useful to limit the number of alerts sent so that an alert is not sent every minute.

If *Exceeds Threshold Limit* is **100** and *Over Timeframe* is **60**, then the threshold condition will only match when this rules other conditions match approximately 100 times over any 60 second period. If *Group Field* is set to "CClientAddr", then the threshold load is grouped by "CClientAddr" value in the event objects. Using the above example, this would mean that the alert would only fire when a specific "CClientAddr" like "**192.168.1.100**" does something over 100 times within 60 seconds. The threshold value for other clients like "**192.168.1.150**" is tracked separately.

Adding Alert Rules

Writing and designing alert rules can be a bit of an art.

Start by finding an event that describes the action that you want to be alerted about. The Event Definitions describes all the event objects and the associated attributes on each object.

Set the *Class* to the event you want to alert about, then add conditions that check the fields to look for the events you are interested in.

As an example, lets say we want to setup an alert for when a specific website is visited by a specific user.

As a *Class* select *HttpRequestEvent* Then as a field add *domain = example.com* and *sessionEvent.username = example_user*.

We want to know if this user visits this website a single time, so we do not want to add a threshold. We want it to log this alert so we want to check *Log* and we want to send an email so we're going to check *Send Email*.

However, in reality when a user visits a website it does a lot of separate HTTP requests to load all of the components. We do not want to receive 20 emails each time a user visits a single page on that website. As we want to check *Limit Send Frequency* to *20* minutes so that we aren't flooded with emails.

There are many other alert rules that are not enabled by default that can provide some common examples.

Triggers

Triggers are similar to Alert rules, however instead of alerting when something interesting happens, trigger rules can "tag" a specific host, device, or user for a specific period of time.

Unlike most rules, all Trigger rules are evaluated even beyond the first matching rule.

This allows the system to keep state on the different hosts on the network, which can serve several purposes. For example, you can tag a specific host/device/user as using a specific application when the application is used.

There are several rules that are included but not enabled to provide some examples.

- **Enable** determine if the alert rule is enabled.
- **Class** is the type of event this rule matches. Selecting the *Class* will determine what *Fields* are available in the conditions.
- **Conditions** lists the fields within the event object to be checked. If all of the conditions match then the rule will match.
- **Enable Thresholds** limits the alert from firing until it reaches a certain frequency threshold.

    - **Exceeds Threshold Limit** is the frequency limit for which this condition will match. If the frequency is greater than this value, then the threshold conditions matches.
    - **Over Timeframe** defines the time range, in seconds, to use to compute the frequency.
    - **Grouping Field** defines how to group thresholds by an attribute field in the events. This field is optional.

- **Action Type** determines the action taken

- **Tag Host** will tag the specified host with the specified tag.
- **Untag Host** will remove the specified tag from the specified host.
- **Tag User** will tag the specified user with the specified tag.
- **Untag User** will remove the specified tag from the specified user.
- **Tag Device** will tag the specified device with the specified tag.
- **Untag Device** will remove the specified tag from the specified device.
- • **Target** identifies the specific host/device/user. If its a single attribute name for example 'cClientAddr' it will look up to 3 layers deep within an object for any attribute named cClientAddr. If it is a fully qualified name like 'sessionEvent.cClientAddr' it will look at that specific attribute within the specified sub-object.
- **Tag Name** specified the string (name) of the tag to be given or removed.
- **Tag Lifetime** specifies the lifetime of the tag when adding a tag. After the lifetime expires the tag will disappear.

Syslog

Syslog sends events via syslog messages to a remote syslog server. To use syslog simply install a syslog *receiver* on another server, then enable syslog and configure as necessary.. Some syslog products are easier to set up than others. Kiwi, a third-party syslog daemon, is a favorite of many users using Windows, while those on *nix can use rsyslog.

- **Host**: The host name or IP address of the Syslog daemon that is authorized to receive syslog messages from the NG Firewall server. Do **not** set the Host to the NG Firewall itself - this will result in the hard drive filling up very quickly and most likely crashing the box.
- **Port**: The UDP port to send syslog messages to the syslog daemon. 514 is the default as this is the default syslog port.
- **Protocol**: The protocol to use to send syslog messages. The default is UDP.

Syslog Rules

WARNING: Syslog can be a very expensive operation. If configured to send all (or most) events it can seriously negatively impact the performance of the server.

Syslog Rules determine which events are sent via syslog.

Unlike most rules, all Syslog rules are evaluated even beyond the first matching rule.

- **Enable** determine if the alert rule is enabled.
- **Class** is the type of event this rule matches. Selecting the *Class* will determine what *Fields* are available in the conditions.
- **Conditions** lists the fields within the event object to be checked. If all of the conditions match then the rule will match.
- **Enable Thresholds** limits the alert from firing until it reaches a certain frequency threshold.
-
  - **Exceeds Threshold Limit** is the frequency limit for which this condition will match. If the frequency is greater than this value, then the threshold conditions matches.
  - **Over Timeframe** defines the time range, in seconds, to use to compute the frequency.
  - **Grouping Field** defines how to group thresholds by an attribute field in the events. This field is optional.
- • **Remote Syslog** determines if the event is sent via syslog.

To send all events via syslog simply create one rule where *Class = All* and no conditions.

To send specific events to a syslog server configure the *Syslog Rules* to send the specific events to the syslog server.

Email Template

You can customize the content of email alerts by editing the Email Template. Items surrounded by the percent symbol represent system variables. You can use these throughout the **Subject** or **Body** of the message. The table below describes each variable.

Variable	Information
System.company	Your company name defined in Branding Manager.
Alert.description	The event description of the associated alert rule.
System.host	The Hostname of your NG Firewall system.
Event.class	The event class of the associated alert rule.
Event.summary	The event summary of the associated alert rule.
Event.values.keyvalue	The extended event details of the associated alert rule.

The preview window shows in real time how your changes to the **Subject** or **Body** will appear in the email message content.

3.5 Local Directory

Local Directory stores a list of users that can be used by the applications. It also supports RADIUS for 802.1x authentication from properly configured wireless network access points.

The RADIUS Server can be enabled to allow WiFi users to authenticate as any user configured in *Local Directory*.
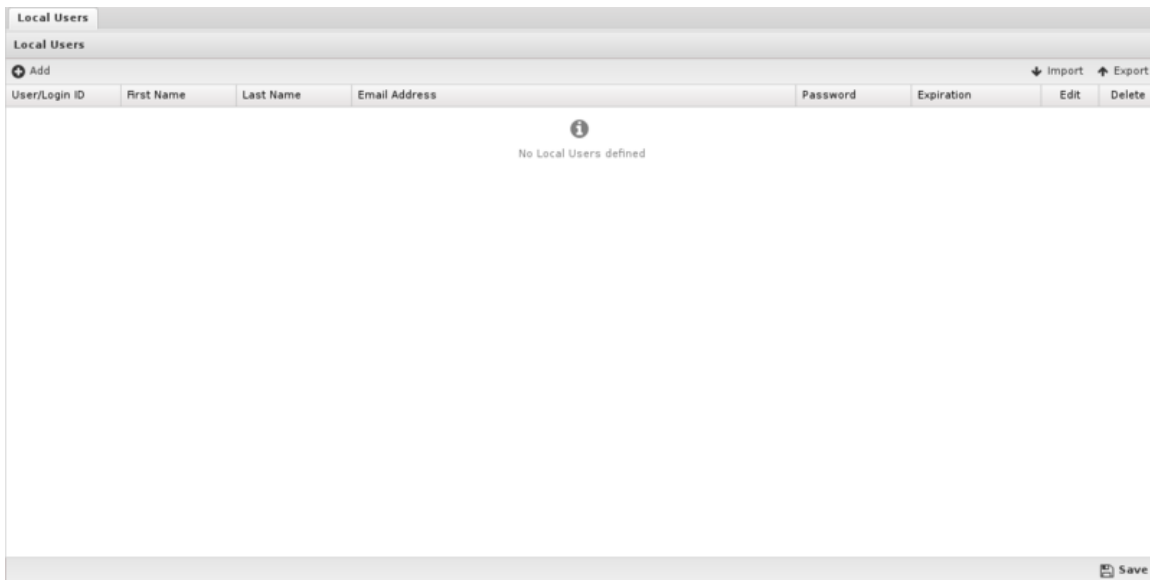
The RADIUS Proxy can be enabled to allow WiFi users to authenticate with credentials that are validated with a configured Active Directory Server.

- Local Users
- RADIUS Server
- RADIUS Proxy
- RADIUS Log

3.5.1 Local Users

Local Users stores a list of users that can be used by the applications.

For example, Captive Portal and OpenVPN can use the local directory to authenticate users.

To add new users click on the Add button. You must supply a username, first name, last name, email address, and password. Only the administrator can set the password for a given user. Users can be imported or exported using the import/export buttons on the upper right.
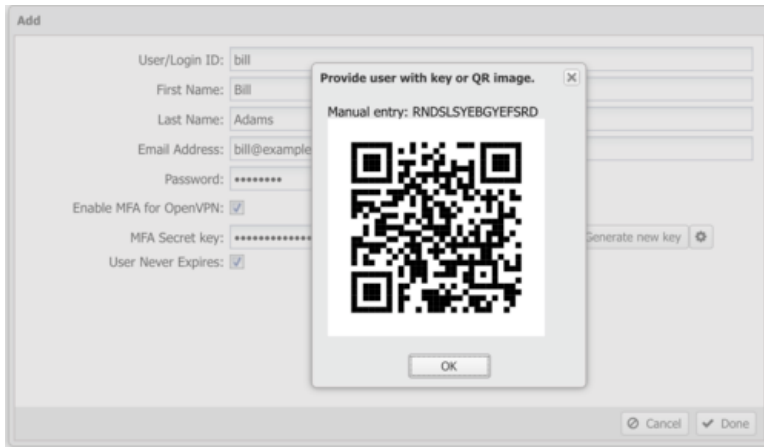
An expiration date can be specified for a user. If the expiration date has passed that user will no longer be authenticated.

To use the Local Directory, configure apps such as Captive Portal and OpenVPN to authenticate against the Local Directory while requiring user authentication.

MFA and OpenVPN

You can enable TOTP based multi factor authentication for OpenVPN client connections. When adding a user, select **Enable MFA for OpenVPN** and click **Generate new key**.


After generating a key, click the gear icon to show the QR code. Use the generated code in any TOTP mobile app such as Google Authenticator. The TOTP app generates a temporary that the user enters into their OpenVPN client. **Note:** You must also enable MFA for client configurations in OpenVPN.

> ⚠️ **Warning:** Typically, when passwords are stored, password hashes are saved and the original cleartext password is forgotten so administrators do not have access to user passwords. However, The passwords for users that are stored in the local directory are stored in cleartext because some applications and features (L2TP) depend on access to the cleartext password. Administrators do have access to cleartext user passwords and caution is advised.

3.5.2   RADIUS Log

Radius Log can be used to view the diagnostic messages generated by the RADIUS server.
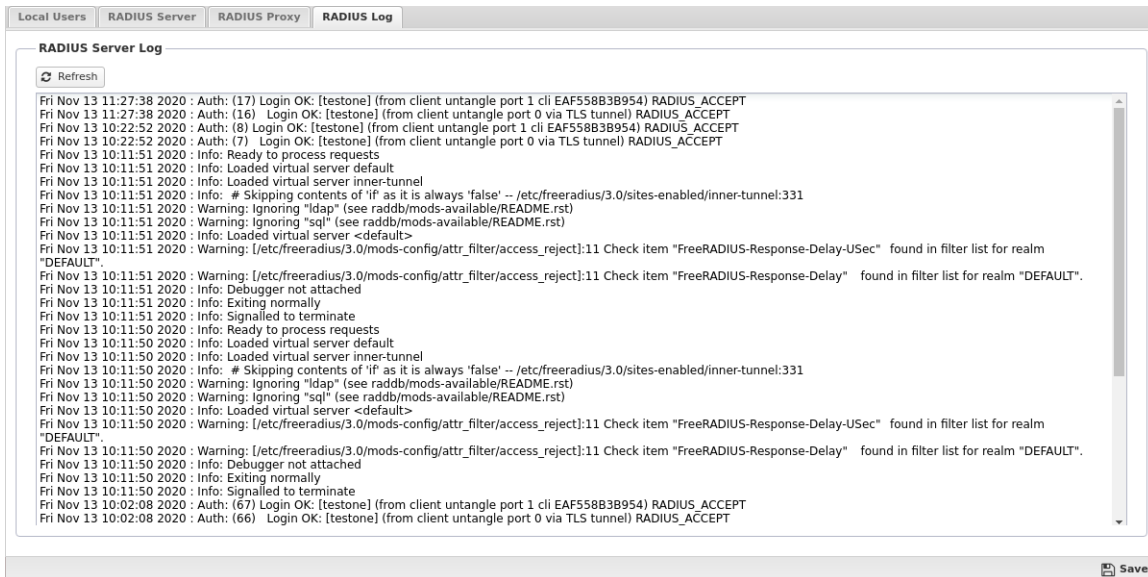
3.5.3   RADIUS Proxy

Radius Proxy is an optional configuration of the RADIUS Server that enables 802.1X authentication against an Active Directory server. Access points configured with WPA/WPA2 Enterprise authentication to the NG Firewall RADIUS Server can enforce login via Active Directory when joining the wireless network.

Prerequisites

The NG Firewall appliance must be able to resolve the fully qualified hostname of your Active Directory Primary Domain Controller. You can test name resolution using the DNS Test of the Troubleshooting utility. If the test fails, you must create a Static DNS Entry in the NG Firewall DNS Server.

Active Directory Server

There are three steps required to configure and verify your configuration with the RADIUS Proxy.

1. Input the Active Directory Server details and click Save to apply and activate the settings.

> **Note:** The **AD Workgroup** should be in upper case.

⊕ Back to Config	▭ **Local Directory**

Local Users	**RADIUS Server**	**RADIUS Proxy**	**RADIUS Log**

The RADIUS Proxy can be enabled to allow wireless clients to authenticate using account credentials stored in an Active Directory server.

☑ Enable Active Directory Proxy

Active Directory Server

AD Server:	adserver.acme.example.local
AD Workgroup:	ACME
AD Domain:	acme.example.local
AD Admin Username:	administrator
AD Admin Password:	••••••••• 👁

2. Click the **Create AD Computer Account** button to register the NG Firewall server with the Active Directory server. If the operation is successful you should see the *distinguishedName*, *whenCreated*, and *whenChanged* fields in the *AD Account Status* field.

Active Directory Status

⟳ Refresh AD Account Status	⚷ Create AD Computer Account

3. Enter a valid username and password in the *Active Directory Test* area and click **Test Authentication**. You should see a message indicating the test was successful.

Active Directory Test

Test Username:	bob
Test Password:	••••••••••••• 👁

⚙ Test Authentication

3.5.4   RADIUS Server

The RADIUS Server enables 802.1x wireless access points to enforce authentication via WPA/WPA2 Enterprise. WPA/WPA2 Enterprise wireless networking provides an optimal level of network authorization by requiring each wireless device to authenticate with unique credentials of an authorized user rather than a shared password. Users can be authenticated against Local Users or Active Directory via the RADIUS Proxy .

Wi-Fi Authentication

To enable support for WPA/WPA2 Enterprise authentication, navigate to the RADIUS Server tab of the Local Directory and select Enable external access point authentication. In the RADIUS Password field assign a strong password.

After you enable the NG Firewall RADIUS server you need to configure your access point to use WPA/WPA2 Enterprise. The following parameters may be necessary to configure WPA/WPA2 Enterprise for your access point.

- RADIUS Server IP address - the IP address of your NG Firewall server on the same LAN segment as your wireless access point.
- RADIUS port number - the NG Firewall RADIUS authentication server listens on port **1812**.
- RADIUS accounting port - the NG Firewall RADIUS accounting server listens on port **1813**. This parameter is optional and may not be supported or configurable on some access points. RADIUS accounting is used by the access point to inform the NG Firewall server about login and logout activities of each authenticated user and their associated device address.
- Shared Secret - the shared secret may also be called a password or key and is used to authorize communication between the access point and the NG Firewall RADIUS server.

Server Certificate

When clients connect to the wireless network, they must install the root certificate of your server. See Certificates. Most devices supporting WPA/WPA2 Enterprise authentication prompt the user to install the certificate when joining the network for the first time.

Access Rules

By default there are two Access Rules to allow access to the RADIUS server from WAN or non-WAN interfaces. The access rules permit UDP protocol to ports **1812** and **1813**. If your access point or domain controller does not belong to a local network then you need to enable the rule Allow RADIUS on WANs.

Note: When upgrading to version **16.2** these rules are not automatically created. You must create them manually to permit access from your access point to NG Firewall.

3.6   System

System contains settings related to the server

- System Reports

    The Reports tab provides a view of all performance reports for the system, including cpu, memory, and disk usage.
- Regional

    The tab configures the region/location specific settings of the NG Firewall server.
- Support

    The tab configures the support settings and allows for rebooting and shutting down the server for support purposes.
- Logs

    The Logs tab configures the number of log files to retain for each log type.
- Backup

    You can export NG Firewall configuration to a local file. This includes all the settings in Config and the settings from the applications.
- Restore

    Restore allows for the restoring of settings from backups created in **Config > System > Backup** or by the Configuration Backup application.
- Protocols

    The protocols tab configures how certain protocol parsing and processing functions.
- Shield

    The shield monitors the session creation rate of the clients creating sessions. Each time a session is processed by NG Firewall the shield calculates the current session creation rate of the client initiating the session.
- System Reports

    The Reports tab provides a view of all performance reports for the system, including cpu, memory, and disk usage.

3.6.1   System Reports

The Reports tab provides a view of all performance reports for the system, including cpu, memory, and disk usage.

Reports

System reports can be accessed via the Reports tab at the top or the Reports tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the Conditions window at the bottom of the page. The data used in the report can be obtained on the Current Data window on the right.

Table 4: Pre-defined Report Queries

Report Entry	Description
CPU Load	he CPU load over time.
Disk Usage	The disk utilization over time.

Report Entry	Description
Memory Usage	The amount of free memory over time.
Swap Usage	The swap utilization over time as a percent of total swap size .
Swap Usage Bytes	The swap utilization over time.
Highest Active Hosts	The highest number of active hosts.
Server Status Events	All system status events.

The tables queried to render these reports:

- Database Schema

Related Topics

- Report Viewer
- Reports

3.6.2   Regional

The tab configures the region/location specific settings of the NG Firewall server.

Current Time

This fields displays the current time on the NG Firewall Server.

Timezone

This is the configured timezone. It is important to have the correct timezone configured to adjust for any time changes that take place throughout the year.

Language

This is the configured language for the NG Firewall server. The administration UI will be displayed in this language, as well as user-visible pages such as the quarantine and block pages. This will not change the language on certain strings like product names and all online services like the account management, help, and store pages.

Regional Formats

The appropriate formatting of numbers and dates may vary depending on your geographical location. While language display settings contain the most appropriate formats for that language, you may wish to override these vales.

- **Use Defaults:** Use the value provided with the language.
- **Override:** Specify different format values for the following fields:

  - **Decimal Separator:** This string is used to separate decimal spaces. For example a period (.) for 1.23.
  - **Thousand Separator** This string is used to separate thousands places. For example a comma (,) for 1,000.
  - **Date Format** This string is used to generate the date display.
  - **Timestamp Format:** This string is used to generate the time display.

  Both Date and Timestamp Formats can use the formatting fields described on the time and date formatting page.

Force Time Sync

This button allows you to force the server time with the internet (via NTP).

Beware if your server time is significantly in the future (hours or days) then force syncing the time may cause issues as the server as the time will go backwards. Threads and processes that are sleeping until a certain calendar date will now awaken at the planned time as the server time has moved significantly backwards. To avoid this it is suggested to reboot after forcing the time to synchronize if the time was significantly off. Also, logs and reports may behave oddly and certain time periods will now occur twice.

3.6.3   Support

The tab configures the support settings and allows for rebooting and shutting down the server for support purposes.

If **Connect to Command Center** is enabled the NG Firewall server will maintain a secure connection with our cloud infrastructure. This channel can be used for centralized management, monitoring, or hot fixes from the cloud.

Note: The NG Firewall server will connect to the Support system (outbound). This does not require you to change any settings on any firewalls in front of the Untangle server to allow for inbound sessions.

If **Allow** secure access to your server for support purposes. is checked then the Edge Threat Management Support team will be able to access your server.

Manual Reboot

This button will reboot the NG Firewall server.

Note: Rebooting should not be done excessively. It will not solve any persistent problems. For more information please refer to A Reboot Fixed It.

Manual Shutdown

This button will power off the NG Firewall server.

Setup Wizard

This button allows you to re-run the Setup Wizard that is automatically launched on a new install.

3.6.4   Logs

The Logs tab configures the number of log files to retain for each log type.

Log Retention

Disk space used by logs: Shows the current usage on the disk occupied by the log files.

For each log type, number of logs to retain: Sets the number of log files to keep for each log type. Choose a low value to limit the amount of space used by logs. The minimum value is *1*.

3.6.5   Backup

You can export NG Firewall configuration to a local file. This includes all the settings in Config and the settings from the applications. It does not include the reporting data, the quarantine data, nor any unique "configuration" like the UID of the server. To export the configuration, go to **Config > System > Backup tab** and click **Backup** to file. For automated configuration backup of configuration and other data, install the Configuration Backup app.
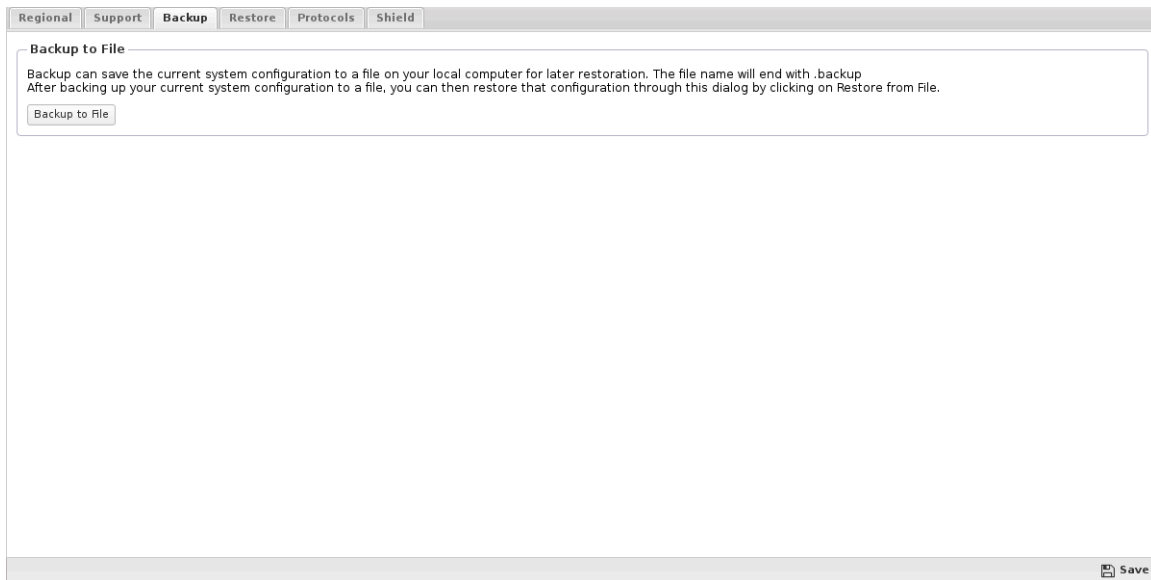
3.6.6   Restore

Restore allows for the restoring of settings from backups created in **Config > System > Backup** or by the Configuration Backup application.
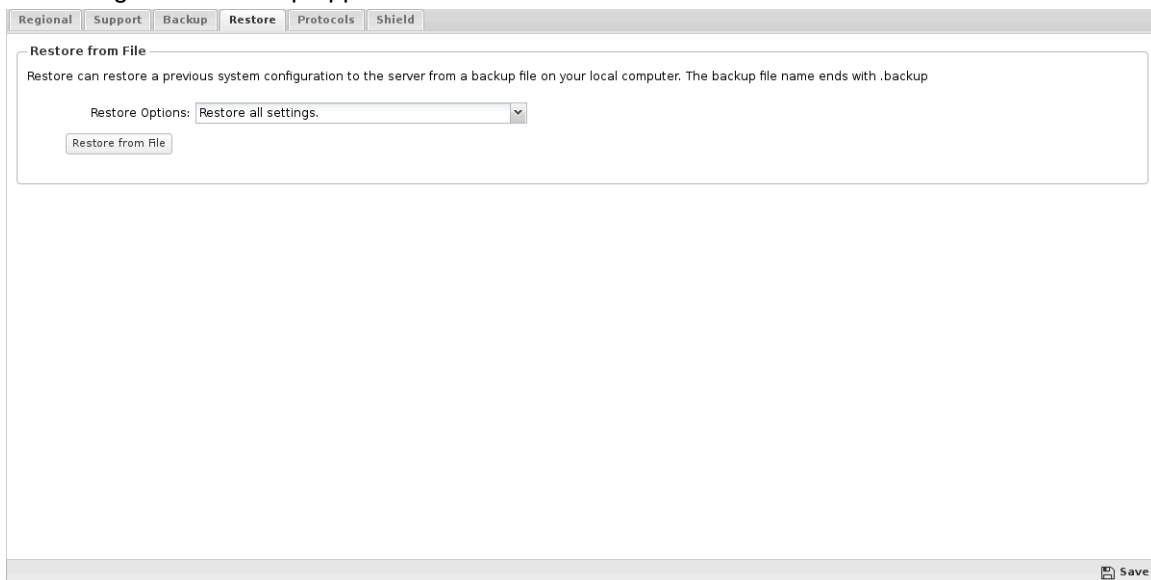
Restore from File

This allows you to upload the restore file.

First, select the Restore Options appropriate for your case.

- Restore all Settings will restore all the settings in the backup file.
- Restore all except keep current network settings will restore all the settings in the backup file except the network settings. The current network settings will be maintained.

The first option is typically used if you are restoring to a previous backup or recovering from a failure.

The second option is useful if you maintain a 'standard configuration' and you want to maintain this standard configuration across multiple servers. In this case all the servers maintain the same settings, but each have unique network settings.

After selecting the Restore Options click browse and select the backup file you wish to restore. After selecting the backup file click Restore from File to begin the restore process.

Restore Process

After starting the restore process the backup file is unpacked and checked.

If the backup file requires certain applications that are not currently on the NG Firewall server, it will ask to download these applications first. After downloading those applications the restore process can be run again.

If the backup file is from an unsupported version it will show an error. It is also suggested to restore a backup file from the same version that the file was created with. For example, if the backup file was created with *NG Firewall 16.2*, then it is suggested to restore on a NG Firewall running *16.2*.

Typically the only supported versions of the restore process will be the current version of NG Firewall and the immediately prior major version. For example, *16.2* will restore *16.2* backups and *16.1* backups, but not *16.0*. (Trivial versions are considered identical to the minor version for restore purposes. For example, *15.1.0, 15.1.1*, and 15.1.2 are all considered 15.1 when restoring backups.)

After the restore process begins the NG Firewall processes will reboot and you will lose connection to the server. After reconnecting to the server you will see the settings and configuration restored from the backup file.

3.6.7   Protocols

The protocols tab configures how certain protocol parsing and processing functions.

⚠️ **Warning:** These settings should not be changed unless instructed to do so by support.

Figure 20: Protocols Tab

The protocols that appear and the visibility of the *Protocols* tab depends on the current applications installed. Many applications use hidden applications dedicated to the processing and handling of important protocols like SMTP and HTTP. These hidden applications can be enabled/disabled in this tab.

If a protocol is disabled, those sessions will still be treated as a binary stream but will not be parsed and unparsed.

HTTP

If enabled, the http-casing HTTP processing application is running. If disabled, all special HTTP processing by the applications stops.

SMTP

If enabled, the smtp-casing SMTP processing application is running. If disabled, all special SMTP processing by the applications stops.

FTP

If enabled, the ftp-casing FTP processing application is running. If disabled, all special FTP processing by the applications stops.

3.6.8   Shield

The shield monitors the session creation rate of the clients creating sessions. Each time a session is processed by NG Firewall the shield calculates the current session creation rate of the client initiating the session. If the session creation rate of the client reaches a level that the shield considers too aggressive the session creation rate of that client is limited to that level.

This process protects the NG Firewall server and also protects the network from Denial of Service (DOS) attacks.

Enable Shield

If checked, the shield is enabled. If unchecked the shield is disabled. Warning: do not disable the shield. Doing so may cause performance and stability issues. This checkbox is provided to allow for troubleshooting. It is never suggested to leave the shield disabled after any troubleshooting steps.

Note, the shield only looks at new session requests, it does not influence or process traffic of existing sessions. It also does not scan bypassed sessions.

Shield Rules

Shield rules are evaluated at session creation time. The Rules documentation describes how rules are processed.

If one of the rules matches, the action from the first matching rules is applied. If no shield rule matches the session is scanned.

If the session is scanned if the current session creation rate is too high, the packet will be dropped. If the current session creation rate is not too high, the current session creation rate is adjusted to account for this new session and the session is allowed.

Reports

The Reports tab provides a view of all reports and events for all traffic handled by Shield.

Reports

This applications reports can be accessed via the Reports tab at the top or the Reports tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the Conditions window at the bottom of the page. The data used in the report can be obtained on the Current Data window on the right.

Table 5: Pre-defined Report Queries

Report Entry	Description
Scanned Sessions	The amount of scanned and blocked sessions over time.
Blocked Sessions	The amount of blocked sessions over time.
Top Blocked Usernames	The number of blocked sessions grouped by username.
Top Blocked Clients	The number of blocked sessions grouped by client.
Top Blocked Ports	The number of blocked sessions grouped by server port.
Top Blocked Servers	The number of blocked sessions grouped by server.
Top Blocked Hostnames	The number of blocked sessions grouped by hostname.
Scanned Session Events	All sessions scanned by Shield.
Blocked Session Events	All sessions blocked by Shield.

The tables queried to render these reports:

- Database Schema

3.6.8.1   Shield Reports

The Reports tab provides a view of all reports and events for all traffic handled by Shield.

Reports

This applications reports can be accessed via the Reports tab at the top or the Reports tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the Conditions window at the bottom of the page. The data used in the report can be obtained on the Current Data window on the right.

Table 6: Pre-defined Report Queries:

Report Entry	Description
Scanned Sessions	The amount of scanned and blocked sessions over time.
Blocked Sessions	The amount of blocked sessions over time.

Report Entry	Description
Top Blocked Usernames	The number of blocked sessions grouped by username.
Top Blocked Clients	The number of blocked sessions grouped by client.
Top Blocked Ports	The number of blocked sessions grouped by server port.
Top Blocked Servers	The number of blocked sessions grouped by server.
Top Blocked Hostnames	The number of blocked sessions grouped by hostname.
Scanned Session Events	All sessions scanned by Shield.
Blocked Session Events	All sessions blocked by Shield.

The tables queried to render these reports:

- Database Schema

Related Topics

- Report Viewer
- Reports

3.7   Email

Email contains all the email-related configuration of the NG Firewall server.

- Outgoing Server
- Safe List
- Quarantine

3.7.1   Safe List

The safe list is a list of email addresses that are considered safe or trusted.

Safe List

Several applications, such as Spam Blocker, Spam Blocker Lite, and Phish Blocker scan SMTP messages.
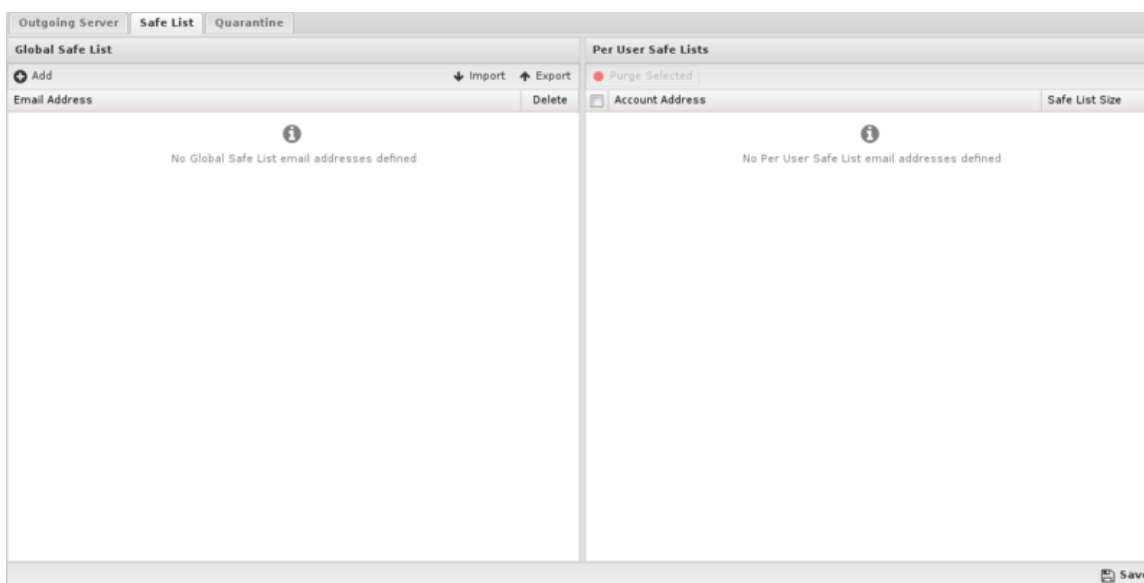
Figure 21: Email Safe List

Sometimes administrators wish to trust email from certain addresses to avoid scanning messages to either save resources or avoid false positives. The safe list provides a convenient location to list safe and trusted email addresses that these applications will check before scanning email.

Note: Virus Blocker and Virus Blocker Lite do not check the safe list because of the low false positive rate.

Global Safe List

This is a global safe list that applies to all email. If an email address is listed, all mail from that address will not be scanned in Spam Blocker Spam Blocker Lite, and Phish Blocker.

Emails can be specified using Glob Matcher syntax so, for example, you can safe list entire domains as "*@example.com."

Per-User Safe List

Each user/email address also has their own safe list. For example, lets assume "user@example.com" has a quarantine that they manage via the quarantine application. In the quarantine application they can add addresses to their own safe list.

For example, user@example.com may add "spammyemailer@chainletters.com" to their safelist. All email from "spammyemailer@chainletters.com" to "user@example.com" will automatically be passed as safelisted, while email from "spammyemailer@chainletters.com" to other users will be scanned as normal.

Per-User safe lists provide a mechanism to deal with false positives that won't effect the overall false negative rate of other user/emails.

Emails can be specified using Glob Matcher syntax so, for example, you can safe list entire domains as "*@example.com." Also, note that a user/email can add "*" to their Per-User safe list to entirely disable spam/phish scanning for emails to them.

Users can edit their own Per-User safe list in the quarantine web application. Administrators can purge users's safe lists in the administration UI.

3.7.2   Server

This tab shows the current information about the Arista server.

Unique ID

The first field shows the Unique ID (UID) of the Arista server. The UID is a 16 alphanumeric code that uniquely identifies this server for licensing and tracking purposes.

Never share the UID of the server.

The UID is generated automatically on install and each server must have a unique UID to function properly.

Cloning servers post-installation will create two servers with identical UIDs which will result in problems and licensing issues.

Server Information

The second field shows the build version and server information.

Build shows the version of the Arista-vm.

Kernel shows the kernel version. The other fields are used by Arista support.

Current "licensed" device count shows the current number of devices in the host table that count as "licensed" devices.

Highest "licensed" device count since reboot shows the highest value of licensed devices seen by this Arista since reboot.

3.7.3   Outgoing Server

This configures how NG Firewall will send email.

The NG Firewall server sends emails for several reasons:

- The Quarantine facility sends users a daily digest of the spams they received.
- The Quarantine allows users to "release" emails from the quarantine.
- The Reports sends daily summary reports to administrators about NG Firewall server activity.

For these functions to work correctly, NG Firewall must be configured correctly to be able to send email in your environment.

Outgoing Email Server

If **Send email** directly is checked then NG Firewall will send email like a regular email server. It does this by looking up the MX DNS record of the recipient domain and sends the message via SMTP to that address. This generally works with no further configuration, however many residential and even commercial ISPs block port *25* to prevent spam and this will prevent NG Firewall from sending email.

If Send email using the specified SMTP Server is checked then NG Firewall will send email using the configured server as an SMTP relay. In order for this to work **the SMTP relay must be configured to allow NG Firewall to relay email.**

- Server Address or Hostname is the IP address or hostname of the SMTP relay.
- Server Port is the port to use to connect to the SMTP relay.
- If Use Authentication is checked then NG Firewall with authenticate with the SMTP relay.

    - Login configures the username to use during SMTP authentication.
    - Password configure the password to use during SMTP authentication.

Email from Address

This is the "**from**" address of all email sent from the NG Firewall server (excluding emails released from the quarantine).

Email Test

This sends a test email from the configured **Email from Address**. If your email settings are correct the specified recipient should receive the test email within a few minutes.

3.7.4 Quarantine

Spam Blocker Spam Blocker Lite, and Phish Blocker sometimes determine an email is spam or phish and the email should be dropped. However, just dropping an email can be dangerous as perhaps it is a "false positive" and is actually an important email. In this case, dropping the email would be very bad.

The **quarantine** action in these applications exists so that important emails don't get lost. The quarantine action will silently send email to the user's quarantine. All the suspected spam/phish email sits in the quarantine and the user is free to review quarantined email to verify nothing important was quarantined.

If something legitimate was quarantined (called a false positive) the user can Release the email to their inbox.

Quarantine Web Application

Each day, users/emails with new email in their quarantine will be sent an Quarantine Digest email with a link to their quarantine. Alternatively, users can request a Quarantine Digest email by accessing https://NGFW_IP:HTTPS_PORT/quarantine/.

After clicking on the **Click her**e to access your spam quarantine link, the user can view the Quarantine web application which allows them to manage their quarantine and their Safe List

Quarantined Messages: The Quarantine Messages tab shows the list of messages currently in the quarantine. To release messages to the inbox simply check the message(s) and click on **Release to Inbox**. To release a message and automatically add the sender to your safe list click on **Release to Inbox & Add Senders to Safe List**. To delete message select the message(s) and click on **Delete**. Note it is not necessary to delete messages, messages will automatically be purged from the quarantine after the configured time elapses.

Safe List: This tab configures your safe (trusted) email addresses. Email from the listed address will not be scanned to determine if they are spam or phish. If a user's email is falsely determined to be spam, their email address can be added to this list to determine it does not happen again.

Forward or Receive Quarantines: Often mailing lists or aliases will receive Quarantine Digests. This is annoying as all users on the list will receive the Quarantine Digest email. To avoid this you can forward the quarantined mail to a another user's quarantine, like the email list's administrator. Email will still be quarantined and released like normal, but the administrator can do it via their quarantine. Forward Quarantined Messages To configures where quarantined messages will be placed. Received Quarantined Messages From shows any other addresses you are receiving quarantined messages from.

Quarantine Settings

The quarantine behavior can be configured via the administration UI in **Config > Email > Quarantine**.

- Maximum Holding Time (days) configures how long email will be held in a quarantine before it is automatically deleted.
- Send Daily Quarantine Digest Emails configures if daily emails will be sent to users with new mail in their quarantine.
- Quarantine Digest Sending Time configures when the daily digests will be sent, if enabled.

User Quarantines: This shows a list of currently existing user quarantines. User quarantines are created dynamically when an email is quarantined for an email address. There is no need to delete quarantines, this will happen automatically when there are no messages. To release or purge (delete) a user's entire quarantine select the appropriate row(s) and click on the **Purge Selected** or **Release Selected** button at the top. To view a user's quarantine, click on the **Show Detail** icon on the appropriate row. This will display a window showing all the existing messages in that user's quarantine. Messages and can purged (deleted) or released by clicking on the message(s) and clicking on the **Purge Selected** or **Release Selected** button at the top.

Quarantinable Addresses: This is a list of emails that will have quarantines automatically created on their behalf. Sometimes you want to assure that quarantine is not an option for some scanned mail. As such you can put "*@mydomain.com" and only "@mydomain.com" email addresses will have quarantines created dynamically. If an email is scanned for another address and the action is quarantine but it is not a quarantinable address it will be marked instead.

> **Note:** this should almost always be a list with one entry of "*". This means all emails will have quarantines created for them if spam/phish is caught for them. This is the default and suggested value. Most of the time this is used to compensate for some other misconfiguration like it scanning email it should not be scanning (like outbound email). Changing this setting is not suggested.

Quarantine Forwards: As discussed above, it is often desirable to have distribution lists or aliases for their quarantined email to an administrator's email quarantine so the entire lists does not receive quarantine digest emails. You can view/add/delete forwards in this table.

> **Example:** you may wish to forward quarantined mail for the distribution list "everyone@mycompany.com" to "itadmin@mycompany.com" so that only "itadmin" will

get messages about spam to the distribution list. "itadmin" can then manage spam to "everyone@mycompany.com" in their own quarantine.

Network Configuration

The most critical configuration in NG Firewall is the proper configuration of your network settings in **Config** > **Network**.

For simple networks, the configuration completed during the Setup Wizard is probably sufficient. However, some networks have multiple WANs, multiple LANs, various subnets, VLANs, VRRP, etc. This describes how networking operates and is configured in NG Firewall.

This section includes the following sections:

- Interfaces
- Hostname
- Services
- Port Forward Rules
- NAT Rules
- Bypass Rules
- Filter Rules
- Routes
- DNS Server
- DHCP Server
- Advanced
- Network Reports
- Troubleshooting

4.1 Interfaces

The Interfaces page configures the network interfaces or the server.

Interfaces Grid

The Interfaces tab shows the current interfaces and the current status and some configuration information.
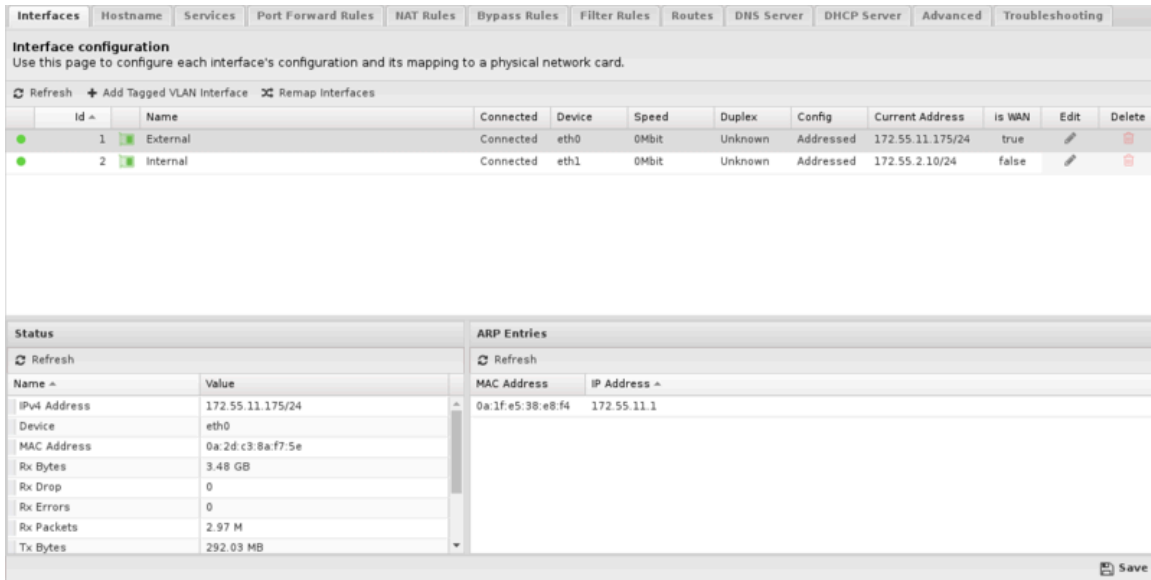
The are several columns along the top of the grid that show the current interface status and configuration. Some are hidden by default.

Columns

Column	Description
Id	The Id is a unique integer primary key of the interface. All configuration of interfaces will refer to Id.
Name	This is a name/description of the interface. It is recommended to choose names representative of their purpose.
Connected	This shows the current "connected" state of the device currently mapped to this interface. This may not display correctly for all network interface cards.
Device	This shows the current network device (physical NIC card or wireless card) mapped to this interface.
Speed	
Duplex	
Config	This shows the current type of configuration for this interface. ADDRESSED, BRIDGED, or DISABLED.
Current Address	This shows the current address if there is one of the interface.
is WAN	This shows true if the interface is configured as a WAN, false otherwise.
Edit	This column shows an edit button to edit the configuration of this interface.
Delete	This column shows an delete button on VLAN Tagged Interfaces to delete the interface. Physical interfaces cannot be deleted, unless their their physical devices have been removed from the system.

There are also several additional options on this page:

- Remap Interfaces

- This utility can be used to change the mapping between physical devices and the corresponding interface configurations. This is useful if you want to use certain physical devices for certain purposes.
- Refresh Device Status

    - This refreshes the "Connected" column in the interfaces grid. To verify your interface mapping plug/unplug one network card at a time and hit *Refresh Device Status* to verify that the expected interface changes the Connected status.
- Add VLAN Tagged Interface

    - This allows for additional of 802.1q VLAN tagged interfaces. For more information read VLANs section in Network Configuration.

Interface Configuration

Clicking the edit button on an interface will open the interface configuration settings for that interface.

An interface can be configured in many ways. Some settings and configuration options are only relevant and/or available in certain configurations. As such, based on an interface's configuration certain options may appear and disappear. For example, when checking 'is WAN' the options available to WAN interfaces will appear. After unchecking 'is WAN' the WAN options will disappear and the options for non-WAN interfaces will appear. Because of this it is suggested to configure your interface from the top of the page downward.

The table below shows the various configuration options and their meanings.

Interface Options

Option	Description
Interface Name	This is a name/description of the interface. It is recommended to choose names representative of their purpose.
is VLAN (802.1q) Interface	This is true if this a tagged VLAN interface. Otherwise this is not shown
Parent Interface	This is the parent interface for this tagged VLAN interface. This is only shown for VLAN interfaces.
802.1q Tag	This is the VLAN tag for this interface. This is only shown for VLAN interfaces.
Is Wireless Interface	This is available if the interface is detected as a wireless (wlan) interface. Otherwise this is not shown.
Config Type	This is the basic configuration type of this interfaces. *Addressed* means this interface has its own address and configuration. *Bridged* means this interface is bridged to another interface. *Disabled* means this interface is entirely disabled.
is WAN Interface	This should be checked if this is a WAN (Wide Area Network) interface. This means it is connected to your ISP or an internet connection. This should be unchecked if this interface is connected to a private/local network.

Wireless Configuration - This section configures the wireless settings for wireless interfaces. This is only shown for wireless interfaces.

Option	Description
SSID	The broadcasted Service Set Identifier (SSID) for the wireless network.
Mode	Access Point (AP) or Client.
Visibility	Select whether to advertise or hide the SSID.
Encryption	Encryption method used for the wireless signal. WPA2 is recommended.
Password	When encryption is enabled, a password will be required to access the network.
Regulatory Country	Choose the country in which this NG Firewall is based. This is required to comply with regulations around Wi-Fi bands & frequencies.
Channel	Choose from the available channels available and 2.4GHz or 5GHz frequencies. The options available here are dependent on your wireless card. **Warning:** Many chips/drivers do not correctly implement "Automatic" (ACS or Automatic Channel Survey) so it may not work depending on your card. **Notice:** Automatic channel selection has been removed from modern builds due to lack of support and usability issues.

IPv4 Options - This section configures the Internet Protocol v4 (IPv4) settings of this interface.

Option	Description
Config Type	This is the IPv4 configuration type. *Static* means this interface has a static IPv4 address. *Auto (DHCP)* means this interface will use DHCP to automatically acquire an address. *PPPoE* means this interface will use PPPoE to acquire an address. This option is only available for WAN interfaces because non-WANs can only be statically configured.
Address	This is the IPv4 static address. It is only shown if Config Type is *Static*
Netmask	This is the IPv4 static netmask. It is only shown if Config Type is *Static*
Gateway	This is the IPv4 static gateway. It is only shown if Config Type is *Static*
Primary DNS	This is the primary DNS used for DNS resolution. It is only shown if Config Type is *Static*
Secondary DNS	This is the secondary DNS used for DNS resolution. It is only shown if Config Type is *Static*
Address Override	If set, this address will be used instead of the one in the offered DHCP lease. It is only shown if Config Type is *Auto (DHCP)*
Netmask Override	If set, this netmask will be used instead of the one in the offered DHCP lease. It is only shown if Config Type is *Auto (DHCP)*

Option	Description
Gateway Override	If set, this gateway will be used instead of the one in the offered DHCP lease. It is only shown if Config Type is *Auto (DHCP)*
Primary DNS Override	If set, this will be used instead of the one in the offered DHCP lease. It is only shown if Config Type is *Auto (DHCP)*
Secondary DNS Override	If set, this will be used instead of the one in the offered DHCP lease. It is only shown if Config Type is *Auto (DHCP)*
Username	This is the PPPoE username. It is only shown in Config Type *PPPoE*
Password	This is the PPPoE password. It is only shown in Config Type *PPPoE*
Use Peer DNS	If checked the server will use the DNS provided by the PPPoE server for DNS resolution. It is only shown in Config Type *PPPoE*
Primary DNS	The primary DNS to be used for DNS resolution. It is only shown in Config Type *PPPoE* and *Use Peer DNS* is unchecked.
Secondary DNS	The secondary DNS to be used for DNS resolution. It is only shown in Config Type *PPPoE* and *Use Peer DNS* is unchecked.
IPv4 Aliases	This is a list of *alias* addresses. This is an additional list of addresses that this interface will have along with their associated netmasks.
IPv4 Options - NAT traffic exiting this interface (and bridged peers)	This option is only available on WAN Interfaces and defaults to checked. If checked all traffic exiting this interface and interfaces bridged to it will be NATd, and all incoming sessions from this interface will be blocked unless they are forwarded via a Port Forward Rules or destined to the local server.
IPv4 Options - NAT traffic coming from this interface (and bridged peers)	This option is only available on non-WAN Interfaces and defaults to unchecked. If checked all traffic coming from this interface and interfaces bridged to it will be NATd, and all incoming sessions to this interface will be blocked unless they are forwarded via a Port Forward Rules.

IPv6 Options - This section configures the Internet Protocol v6 (IPv6) settings of this interface.

Option	Description
Config Type	This is the IPv6 configuration type. *Disabled* means the interface has no IPv6 configuration. *Static* means this interface has a static IPv6 address. *Auto (SLAAC/RA)* means this interface will use SLAAC to automatically acquired an address. This option is only available for WAN interfaces because non-WANs can only be statically configured.

Address	This is the IPv6 static address. Blank is allowed and means no IPv6 address will be given. It is only shown if Config Type is *Static*
Prefix	This is the IPv6 static prefix. It is only shown if Config Type is *Static*
Gateway	This is the IPv6 static gateway. It is only shown if Config Type is *Static*
Primary DNS	This is the primary DNS used for DNS resolution. It is only shown if Config Type is *Static*
Secondary DNS	This is the secondary DNS used for DNS resolution. It is only shown if Config Type is *Static*
IPv6 Aliases	This is a list of *alias* addressed. This is an additional list of addresses that this interface will have along with their associated netmasks. This is only available on non-WAN interfaces.
IPv6 Options - Send Router Advertisements	If checked route advertisements are sent on this interface. This is only available on non-WAN interfaces.

DHCP Configuration (server) - This configures the DHCP serving options on this interface. DHCP Serving is only available on *Addressed* non-WAN interfaces.

Option	Description
Server	If selected, DHCP will be served to this interface so that machines can automatically acquire addresses.
Range Start	The start of the DHCP range.
Range end	The end of the DHCP range.
Lease duration	The duration of the provided DHCP leases in seconds.
Gateway Override	If set, this value will be provided as the gateway in the DHCP leases. Otherwise the static IPv4 address of this interface will be provided as the gateway.
Netmask Override	If set, this value will be provided as the netmask in the DHCP leases. Otherwise the static IPv4 netmask of this interface will be provided as the netmask.
DNS Override	If set, this value will be provided as the DNS in the DHCP leases. Otherwise the static IPv4 address of this interface will be provided as the DNS. A single IPv4 address or a comma-separated list of IPv4 addresses is accepted.

DHCP Options	This is a list of DHCP options for dnsmasq. **WARNING:** this option is for advanced users. The specified DHCP options will be used on this interface. For example, to specify an NTP server use enabled = true, description = "time server", and value = "42,192.168.1.2". For multiple DNS override servers specify enabled = true, description = "DNS", and value = "6,192.168.1.1,192.168.1.2". The value must be specified in a valid dnsmasq format as described in the dnsmasq documentation

DHCP Configuration (relay) - This configures DHCP relay on this interface.

Option	Description
Relay	If selected, DHCP requests received on this interface will be forwarded to a specified DHCP server.
Relay Host Address	The IP address of the relay host server.

Redundancy (VRRP) Configuration - This configures the VRRP redundancy options for this interface. VRRP is only available on statically assigned interfaces. VRRP documentation.

Option	Description
Enable VRRP	If checked, VRRP is enabled on this interface.
VRRP ID	The VRRP (group) ID of this server. Must match the VRRP ID of peers, but must be unique on the server.
VRRP Priority	The VRRP Priority of this server. Higher value is a higher priority. (1-255)
VRRP Aliases	The list of VRRP Virtual Addresses. This list should be the same on all VRRP peers.

Interface Status

The *status* button on the interface brings up a window showing some of the statistics about the interface. This includes statistics, the ARP table, and the connected clients if its a wireless interface.

4.2 Hostname

The tab configures the hostname and related settings of the NG Firewall server.

Hostname

- Hostname

  - This is the name given to the NG Firewall server, such as "NGFW", "firewall", and so on.
- Domain

  - This is the domain name of the NG Firewall server. If your company uses "mycompany.com" you will likely want to use mycompany.com.

The Fully Qualified Domain Name (FQDN) for the NG Firewall server is *Hostname* + *Domain*. So Hostname = "NGFW" and Domain = "mycompany.com" means the FQDN for NG Firewall is *ngfw.mycompany.com*. If you have publicly available services like VPN and/or spam quarantines you should make sure that ngfw.mycompany.com resolves in DNS to the/a public IP of the NG Firewall server.

Dynamic DNS Service Configuration

Several Dynamic DNS services are available to help those with dynamic public IPs. Some ISPs and areas only offer dynamic IPs which can be problematic for networks with remote users wanting to access services. You can not remote users access the server/network by the public IP because it can change at any time.

These services exists to automatically update the public DNS entry when your DHCP address changes. This allows you to refer remote users to a FQDN such as "firewall.mycompany.com" and then automatically update the DNS resolution of "firewall.mycompany.com" to your public IP when it changes.

- Enabled

  - If enabled a Dynamic DNS server will be used to update DNS resolution of the FQDN.
- Service

  - The dropdown shows the supported services. Choose the service you wish to use.
- Username

  - The username to use of the service.
- Password

  - The password of the account of the service.
- Hostname(s)

- The hostname to update with NG Firewall's public IP address. Specify a single FQDN or multiple FQDNs separated by commas.

4.3 Services

The Services tab configures the local services available on the Arista Server. The Arista Server hosts an HTTPS and HTTP server (apache) that hosts services.

Services like administration, spam quarantine, reports, blockpages, etc

- HTTPS Port: TCP Traffic to the primary address of each interface on the HTTPS Port will be forwarded to the local web server to provide services. Often if you only have one public IP (*1.2.3.4*) and wish to port forward the HTTPS port to an internal service like webmail to an internal machine you will need to move the HTTPS Port to another port, like *4343*, so that port can be forwarded to your internal machine. In this case the administration and other HTTPS services will be accessible at https://1.2.3.4:4343/ and *443* can be forwarded.
- HTTP Port: TCP Traffic from non-WANs to the primary address of each non-WAN interface on the HTTP Port will be forwarded to the local web service to provide services. In the rare event that those ports are required for port forwarding the default HTTP Port can be changed to another port, like *8080*.

> **Note:** The configured HTTPS port will be forwarded to the local apache process listening on *port 443*. The configured HTTP port will be forwarded to the local apache process listening on *port 80*. The *Access Rules* will evaluate these sessions post-redirection where destination port has already been altered.

4.4 Port Forward Rules

Port forwarding is a technique of rewriting the destination address and destination port of packets to send them to a new location.

Port Forward Rules

Port forwarding has many uses. The most common use is on networks doing NAT where internal servers have private addresses (ie *192.168.1.100*), port forwarding allows the forwarding of traffic to the NG Firewall server's public IP to an internal host.

For example, if the email server is behind NG Firewall with a private address (*192.168.1.100*) and NG Firewall has one public IP (*1.2.3.4*) which all hosts are "sharing" via NAT to reach the internet, then port forwarding can be used to forward TCP traffic to *1.2.3.4*, *port 25* (SMTP) to *192.168.1.100*, *port 25*.

Port Forwards Rules work like all rules in NG Firewall which is described in the Rules documentation. Rules are evaluated in order on all new sessions. The destination will be rewritten to the *New Destination* and the *New Port* of the first matching rule. If not rule matches, then no changes are made to the session.

If you are having issues with port forwards, it is suggested to follow the Port Forward Troubleshooting Guide.

There are two types port forward rules. To add a *Simple Rule* click **Add Simple Rule.** To add a regular rule as described in the Rules documentation click **Add**.

Simple Rules

Simple rules allow for most use cases and avoid extra configuration, which can lead to non-functional forwards. It is suggested to always use simple port forward rules where possible.

Enable Port Forward Rule: ☑

Description: My Port Forward Rule

Forward the following traffic:

Protocol: TCP ▾

Port: SMTP (25) ▾

Traffic matching the above description destined to any Untangle IP will be forwarded to the new location:

New Destination: 192.168.1.100

[Switch to Advanced]

- Enable Port Forward Rule

  - If checked, the rule is enabled. If unchecked the rule has no effect and is disabled.
- Description

  - A description of this rule. This is just for documentation.
- Protocol

  - If "TCP" then only TCP traffic to the specified port is forwarded. If "UDP" then only UDP traffic to the specified port is forwarded. If "TCP & UDP" both TCP and UDP to the specified port is forwarded.
- Port

  - This is a list of commonly forwarded ports. Use commas to separate multiple ports, or a dash to denote a range. To specify an arbitrary port, select **Other** and specify the port.
- New Destination

  - The new destination of the session after the port forward. Typically this is the internal machine like *192.168.1.100*.

Expert Rules

There are cases where a more complex rule is desired, such as:

- forward non TCP/UDP protocols

  - Simple Rules can only forward TCP and UDP
- changing the new destination port

  - Simple Rules only rewrite the destination address. If you also want to change the destination port, expert rules are necessary. For example, forwarding TCP to *1.2.3.4* port *8080* to *192.168.1.100* port *80*.
- further limiting what traffic is forwarded

  - Expert Rules allow for additional conditions to limit when a port forward will match. For example, you can limit forwarded traffic to traffic from a specific interface, specific IP range, etc.
- only to a certain public IP (not all NG Firewall IPs)

  - Simple Rules match all *Destined Local* traffic as described in Rules. If your NG Firewall has multiple public IPs and you only want to forward traffic to one of them you can use Expert Rules.

⚠ **Important:** The reason simple rules are encouraged is that users tend to misconfigure the port forward rule conditions. Adding **MORE** conditions means your port forward rule will match **LESS** traffic. For example, you might specify to only forward traffic with "Source Interface

= External" and then discover that traffic to a public IP from the internal interface is NOT forwarded because its source interface is NOT external.

A port forward

- Enable Port Forward Rule

  - If checked, the rule is enabled. If unchecked the rule has no effect and is disabled.
- Description

  - A description of this rule. This is just for documentation.
- Conditions

  - The conditions describing which sessions will match.
- New Destination

  - The new destination of the session after the port forward. Typically this is the internal machine like *192.168.1.100*.
- New Port

  - Optional. If blank, the new destination port will remain unchanged. If specified, all matching sessions will be rewritten to the new destination port.

Reservations

At the bottom of the Port Forward Rules tab there is a list of *reserved* ports that can not be forwarded because they are currently in use for NG Firewall services. These services can be moved to different ports in the Services tab if these ports are required for port forwarding.

4.5   NAT Rules

NAT rules allow the rewriting of the source address of traffic.

Enable NAT Rule: ☑
Description: NAT mail server traffic to 1.2.3.5

If all of the following conditions are met:

	Type			Value	Delete
⊕ Add					
	Source Address ▾		is ▾	192.168.1.100	✕

Perform the following action(s):

NAT Type: Custom ▾
New Source: 1.2.3.5

Typically NAT is used so that machines on a private subnet (*10.*.*.**, *192.168.*.**, etc) can share a single public IP address. To do this when a private machine (say *192.168.1.100*) makes a connection to a public server (say google.com) the NG Firewall server rewrites the source address to the public IP address of NG Firewall (say *1.2.3.4*) on the way out. When return traffic in that session returns to *1.2.3.4* it is rewritten back to the internal address, *192.168.1.100*, and forwarded back to the internal server.

NAT

By default, "NAT traffic exiting this interface (and bridged peers)" is checked on WAN Interfaces which enables the NATing of all sessions exiting that WAN interface with the source address of the primary IP of that interface. In other words, all sessions leaving the External interface will use the External interface's primary IP.

There is also another options called "NAT traffic coming from this interface (and bridged peers)" on non-WAN Interfaces. If checked, all traffic coming from the interface will be NATed using the primary IP address of its destination interface.

By default, since only "NAT traffic exiting this interface (and bridged peers)" is checked, NAT is only done on traffic that exits a WAN interface. This means traffic between internal networks will be un-NATed and each can reach each other using the private addresses. If this is not desired, NAT can be done as traffic comes from a non-WAN by checking "NAT traffic coming from this interface (and bridged peers)". This means NAT takes place between local interfaces and no traffic will flow between separate internal networks without explicit port forwards.

> **Note:** Checking "NAT traffic exiting this interface (and bridged peers)" adds an implicit NAT rule to NAT all traffic exiting that interface to *Auto.* It also adds an implicit Filter Rules to block all from that interface that is not to the local server and is not explicitly port forwarded.

> **Note:** hecking "NAT traffic coming from this interface (and bridged peers)" adds an implicit NAT rule to NAT all traffic coming from that interface to *Auto.* It also adds an implicit Filter Rules to block all to that interface that is not to the local server and is not explicitly port forwarded.

NAT Rules

Occasionally additional rules are necessary for more complex NAT setups.

For example, lets assume you have two public IPs *1.2.3.4* and *1.2.3.5*. By default you want all traffic to be NATd to the primary address *1.2.3.4*, but you want your mail server (*192.168.1.100*) to send mail from *1.2.3.5*. To do this you need to add a NAT rule saying that traffic from the mail server should be NATd to *1.2.3.5*. To do so add a rule with Source Address = *192.168.1.100*, where NAT Type = 'Custom' and New Source = '*1.2.3.5*'.

Another common scenario is setting up 1:1 NAT, using a paired Port Forward Rules and a NAT Rule.

A NAT rule

NAT Rules contain several components.

- Enable

  - If checked, the rule is enabled. If unchecked the rule has no effect and is disabled.
- Description

  - A description of this rule. This is just for documentation.
- Conditions

  - The conditions describing which sessions will match. As documented in Condition List of Rules documentation.
- NAT Type

  - *Auto* or *Custom*. *Auto* means the session will be NATd to the primary address of the interface which the session exits. *Custom* allows you to specify a specific IP.
- New Source

  - If *Custom* is selected, specify the IP address that the session will be NATd to.

Like all Rules, the NAT rules are evaluated in order. The session will be NATd according to the first matching rule. If not rule matches, the session will be NATd according to the checkboxes in the Interfaces settings. If no rules match and all NAT options in Interfaces are disabled or do not match the session in question, the session is not NATd and it is sent with the original source address.

4.5.1   1:1 NAT

This section will help you set up 1:1 NAT.

What is 1:1 NAT

1:1 Network Address Translation (NAT) is a mode of NAT that maps one internal address to one external address. For example, if a network has an internal servers at *192.168.1.10*, 1:1 NAT can map *192.168.1.10* to *1.2.3.4* where *1.2.3.4* is an additional external IP address provided by your ISP.

How do I Setup 1:1 NAT?

You need to do three things: first, setup an external IP Address Alias; second, map inbound traffic destined for the external address so it is redirected to the correct internal machine; third, map outbound traffic from the internal machine out the correct external address. In this example, we'll assume you're trying to setup 1:1 NAT for *192.168.1.10* to *1.2.3.4* - you will need to be in advanced mode to configure 1:1 NAT.

Create an IP Address Alias on the WAN interface for *1.2.3.4* with the appropriate netmask provided by your ISP and save it. This can be done at **Config** > **Network** > **Interfaces** on the specific interface and tells Arista to take ownership of that IP.

Create a port forward for inbound sessions and save it. This rule will cause all inbound sessions destined for *1.2.3.4* to be forwarded to *192.168.1.10*: Destination Address: *1.2.3.4* New Destination: *192.168.1.10*.

Create a nat rule for outbound sessions - this causes all outbound sessions from *192.168.1.10* to be NATd to *1.2.3.4*. This is done at **Config** > **Network** > **NAT Rules**: Source Address: *192.168.1.10* NAT Type: Custom New Source: *1.2.3.4* Once this is configured and saved, your 1:1 NAT setup is complete.

How do I Verify 1:1 NAT is Working?

You can check outbound traffic by going to your internal server and visiting whatismyip.com. You can check inbound traffic by testing your port forward. For example, if your internal server is running a web server then visit http://1.2.3.4/ from outside the network - it should load the web server on *192.168.1.10*.

4.6   Bypass Rules

Bypass Rules work like other Rules. They are evaluated in order.

	Rule Id	Enable	Description	Conditions	Bypass	Edit	Delete
⊕	1	☐	Bypass DNS Sessions	Destination Port ⇒ 53	☑	✏	🗑
⊕	2	☑	Bypass VoIP (SIP) Sessions	Destination Port ⇒ 5060	☑	✏	🗑
⊕	3	☑	Bypass VoIP (IAX2) Sessions	Destination Port ⇒ 4569	☑	✏	🗑
⊕	4	☑	Bypass PPTP Sessions	Destination Port ⇒ 1723 · Protocol ⇒ TCP	☑	✏	🗑

Interfaces | Hostname | Services | Port Forward Rules | NAT Rules | **Bypass Rules** | Filter Rules | Routes | DNS Server | DHCP Server | Advanced | Troubleshooting

Bypass Rules control what traffic is scanned by the applications. Bypassed traffic skips application processing. The rules are evaluated in order. Sessions that meet no rule are not bypassed

⊕ Add   ↓ Import   ↑ Export

💾 Save

NG Firewall's applications run in what we call the "Arista VM" - the Arista Virtual Machine. During this process, both UDP and TCP streams are endpointed and their streams are reconstructed at layer 7 (the application layer). The data stream then flows through the applications, and if passed eventually the data is put back into new packets and sent on its way.
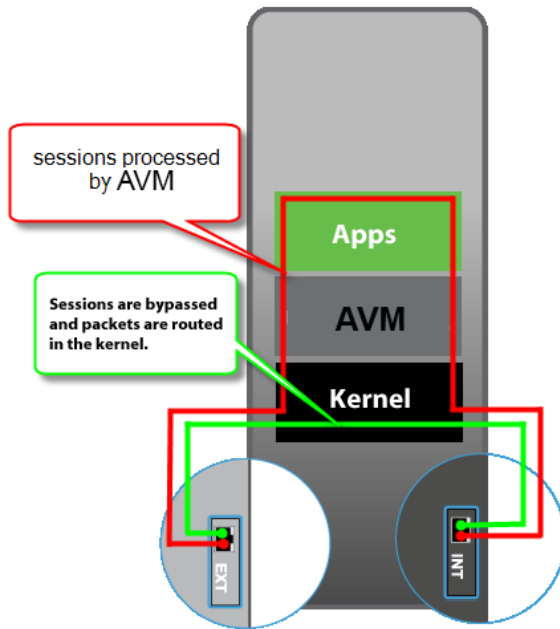
Unlike most proxy firewalls, by default NG Firewall processes almost all ports of both UDP and TCP at the application layer. In some cases, it may be ideal to "bypass" traffic so that it is not subject to scanning. As shown in the image on the right *bypassed* traffic will skip all of the NG Firewall VM layer 7 processing and all the applications.

Sometimes it is ideal to bypass traffic for performance reasons. It can either be traffic that you are not interested in scanning and wish to save server resources, or traffic that is extremely sensitive to scanning, like VoIP.

Sometimes it is also necessary to bypass traffic that the application layer processing interferes with.

Bypassed sessions are still routed, NATd, and filtered identically to all other sessions, as defined by your network configuration. The only difference is that bypassed sessions are not processed at layer 7 so their traffic "bypasses" the applications.

Bypass Rules

Bypass Rules work like other Rules. They are evaluated in order. The action from the first matching rule is taken.

As an example, lets say you have a backup server at *1.2.3.4* and you don't want traffic to that backup server scanned or interfered with in anyway. To bypass it create a rule with Destination Address = *1.2.3.4* and action = "Bypass"

Bypass Rules contain several components.

- Enable

  - If checked, the rule is enabled. If unchecked the rule has no effect and is disabled.
- Description

  - A description of this rule. This is just for documentation.
- Conditions

  - The conditions describing which sessions will match. As documented in the Condition list in Rules section.
- Action

  - *Bypass* or *Process*. *Bypass* means the traffic will be bypassed. *Process* means the traffic will be processed by the AVM and the apps.

Like all Rules, the Bypass rules are evaluated in order. The session will be processed or bypassed according to the first matching rule. If no bypass rule matches, the session will be processed.

Common Uses

There are several scenarios which it usually makes sense to bypass traffic.

On large networks with servers that might be very busy it usually makes sense to bypass traffic that need not be scanned. This all depends on why you are using NG Firewall on your network.

Check in Reports under "System" and look at the "Top Destination Ports." Occasionally you'll see some bizarre port with millions of sessions, like syslog. Often these ports can be bypassed if you don't care to scan them.

Often it makes sense to bypass port **53** from your internal DNS server so you can guarantee that NG Firewall will not interfere with your DNS server's resolution process. This is critical if NG Firewall is using this server for DNS resolution itself.

If you are using NG Firewall for just Web Filter then you can basically bypass all of UDP and save lots of processing time, although obviously this isn't a good idea if you are using Bandwidth Control as then it would not be able to shape UDP.

Bypassing can sometimes be useful for troubleshooting. If you are having issue with some traffic you can bypass it to see if it helps. If it does help then revert back to processing and disable the apps one at a time to see if its one of the applications interfering with the traffic in question.

4.7   Filter Rules

Filter rules are kernel-level iptables (Layer 3) "filter" rules. Filter rules apply to sessions transiting THROUGH the Arista server. By default this ruleset is blank. Filter Rules are useful for blocking traffic going through the Arista server.

- Enable Filter Rule

    - If checked, the rule is enabled. If unchecked the rule has no effect and is disabled.
- IPv6

    - If checked, the filter rule will also be active with IPv6 addressing.
- Description

    - A description of this rule. This is just for documentation.
- Conditions

    - The conditions describing which sessions will match. As documented in the *Condition_List* section of Rules
- Action

    - *Block* or *Pass*. *Block* means the session dropped silently. *Pass* means the session will be passed.

The rules are evaluated in order on all new sessions going through the Arista server as described in the Rules documentation. The action from the first matching rule is taken, if no rule matches the session is passed. All passed sessions are still subject to processing in the Apps.

Why use Filter Rules

There are also block/pass rules available in the Firewall app. There are several key differences that determine when its appropriate to use a Filter Rule vs a Firewall Rule.

- Filter Rules still apply to bypassed traffic. The Firewall does not see bypassed traffic. This means if you want to block anything that's bypassed you should use the Filter Rule.
- Filter Rules apply to all protocols while Firewall only sees TCP and UDP. If you want to block IP protocols other than TCP and UDP, you should use Filter Rules.
- Firewall Rules have more application-layer conditions available like Client has exceeded Quota and HTTP: Client User OS. If you need the application-layer conditions you should use Firewall.
- Firewall Rules are evaluated in the Firewall app, so it can be used in policies setup in Policy Manager.

4.8   Routes

Arista routes **all** traffic according to *its* routing table. As such it is absolutely critical to configure your Arista server with a complete routing table.

If Arista does not have a complete routing table, it will not be able to reach hosts behind Arista and will not properly route return traffic back to them and they will be offline.

Figure 22: Network Routes Example

A route

- Description: A description of this route. This is just for documentation.
- Network: This is the IP/network for this route. Upon saving any bits past the prefix length or netmask will be zeroed out as they are irrelevant and are not accepted in a route.
- Netmask/Prefix: This is the netmask (or prefix) for this route.
- Next Hop: This should either be a currently reachable local IP address, or an interface chosen from the dropdown. If an IP address is specified all traffic to this network will be routed to that IP address. If an interface is specified all traffic to this network will be routed locally on that IP address using ARP to resolve those hosts.

> **Tip:** It is easy to test routes using the Ping Test in **Config** > **Network** > **Troubleshooting** > **Ping Test**. If Arista can ping a host on the network in question it is likely the route is correct. If Arista can not ping a host on the network in question the route is probably incorrect and those hosts will not be online.

Common Uses

Some network have subnets that exists behind other internal routers. For example, lets say my Internal interface is *192.168.1.1/24*. There is also a *192.168.2.0/24* network that exists behind another router at *192.168.1.5*. Without this route that entire network would be offline because its return traffic would go to the wrong place (the default gateway). In this case I need specify that the *192.168.2*.* network is reachable through *192.168.1.5*. This would look as follows:

Some networks also run multiple subnets on the same switch infrastructure internally. If these are "untagged" networks/VLANs, you simply need to add a route. For example, lets say my Internal interface is *192.168.1.1/24*. There is also a *10.0.0.0/8* network on this interface. Without a route these hosts would be offline because its traffic would be routed to the wrong place (the default gateway). In this case I need to specify that *10.*.*.** hosts are reachable on the Internal interface directly. This would look as follows:

4.9   DNS Server

The DNS Server settings configure the DNS server running on the NG Firewall server.

These settings do NOT affect any DNS traffic passing through NG Firewall, only DNS traffic to the NG Firewall server.

It is not required to use the DNS server on NG Firewall, however it is often desired on small networks because the NG Firewall server will cache DNS for the entire network. If NG Firewall is configured 'as a router' where it is providing DHCP to clients on the internal network, the default is to provide the NG Firewall server as the DNS server.

Static DNS Entries

Static DNS Entries are entries that will always resolve to the address provided. Often this is useful for servers hosted internally. For example, if your mail server is local you can add a static entry for

mail.mycompany.com to its internal IP (like **192.168.1.20**). This means machines using NG Firewall for DNS will resolve this hostname to the internal IP and communicate with it directly.

Domain DNS Servers

Often certain domain need to be resolved using certain DNS servers instead of the DNS servers configured on the WAN interfaces. For example you may want all queries to "*.mycompany.local" to go to the local DNS server for resolution. *Domain DNS Servers* allow you to specify that all queries matching *domain* go to the specified server. For example, if all *.example.com queries should go to **192.168.1.20**, then you can add an entry for *Domain* = *example.com* with *Local Server* = **192.168.1.20**.

In this scenario, the NG Firewall and all those using the NG Firewall for DNS resolution will have the matching queries resolved through the specified server. For example, If someone using the NG Firewall server for DNS resolves aaa.example.com this DNS query will be forwarded to **192.168.1.20** instead of NG Firewall's upstream DNS servers configured in the WAN interface settings.

This can also be used to tell NGFW how to do reverse DNS lookups using *in-addr.arpa* as the domain. For example if you wish **172.16.*.*** reverse DNS queries to go to **192.168.1.10**, then set the Domain of "**16.172.in-addr.arpa**" and the Local Server of "**192.168.1.10**". If you wish for **10.*.*.*** reverse DNS queries to go to "**1.2.3.4**" then set the Domain to "**10.in-addr.arpa**" and the Local Server of "**1.2.3.4**".

4.10   DHCP Server

These settings configured the settings of the DHCP Server running on the NG Firewall server.

DHCP Server

> **Note:** The DHCP configuration for each interface is handled in the configuration of that interface in **Config > Network > Interfaces**. This page handles the global DHCP configuration.

Static DHCP Entries

This table contains any *static* DHCP leases. Entries in this table will always be given the same DHCP lease with the configured address. For example if *MAC Address* = *aa:bb:cc:00:11:22' and Address* = **192.168.1.100** *then when the machine with **aa:bb:cc:00:11:22** request a DHCP lease it will always be given **192.168.1.100**.

Current DHCP Leases

This shows the current table of active DHCP leases and their expiration time.

Custom dnsmasq Options

This textfield holds any custom dnsmasq options. **This is for advanced users. Misconfiguration of this field will result in improper functioning of the NG Firewall server.** Any text in this field will be appended to the dnsmasq.conf.

DNSMasq is the server NG Firewall uses to provide DNS & DHCP services. DNSMasq documentation describes the options available.

4.11   Advanced

4.11.1   Options

Options contains some global networking options.

- Enable SIP NAT Helper
  - This enables the kernel SIP NAT fixup. Most SIP solutions handle NAT on their own, but sometimes rewriting of address inside SIP by the NAT device is necessary. Enabling this will enable bypassed SIP sessions to be rewritten in the kernel. Default is off.
- Send ICMP Redirects
  - ICMP Redirects are used to alert machines if a shorter route is available. Default is on.
- Enable Spanning Tree Protocol (STP) on Bridges
  - This enables Spanning Tree Protocol (STP) on bridges which is a protocol used to help detect loops and avoid packet storms in this case. Given that a bridge loop is fatal configuration, this option is off by default so the fatal configuration will fail quickly. It is NOT suggested to rely on STP to stop bridge loops.
- Strict ARP Mode
  - If enabled ARP replies will only go out to for network requests where the request source matches the expected configuration. This helps avoid ARP flux with complicated networks. Strict mode means **arp_ignore** = **1**, **arp_announce** = **2**. Loose mode means **arp_ignore** = **0**, **arp_announce** = **0**. More documentation about **arp_ignore** and **arp_announce** can be found here.
- DHCP Authoritative
  - If enabled, all DHCP serving is authoritative. Default is on. DHCP Authoritative is documented here.
- Block new sessions during network configuration
  - If enabled, all sessions will be blocked (dropped) when network settings changes are applied. This will provide increased security for router mode deployments and is not recommended for bridged mode deployments. The default setting is disabled.
- Log bypassed sessions
  - If enabled, bypassed sessions will be logged to the sessions table
- Log outbound local sessions
  - If enabled, bypassed sessions created by the Arista server itself will be logged to the sessions table
- Log inbound local sessions
  - If enabled, bypassed sessions to the Arista server itself will be logged to the sessions table

- Log blocked sessions

    - If enabled, all sessions blocked by filter rules or NAT or the shield will be logged to the sessions table.

4.11.2 QoS

About QoS

Quality of Service (QoS) is a mechanism to ensure high-quality performance to latency- and bandwidth-sensitive applications. It allows for the prioritization and differential treatment of traffic based on rules. Most often this is used to improve the performance of latency and bandwidth sensitive applications and traffic (like VoIP) at the cost of less important traffic such as peer-to-peer. QoS can greatly improve the performance of the network traffic and important protocols, especially when the upload or download bandwidth is saturated. However QoS can also be detrimental to network performance if configured incorrectly. **It is advised to read this section in its entirety before enabling QoS.**

QoS settings can be found at **Config** > **Network** > **Advanced** > **QoS**.

The Seven Priorities

The seven priorities in the default configuration can be thought of as two sets - the top four priorities: *Very High, High, Medium*, and *Low* can all consume all available bandwidth if no higher priority class wishes to use it. Use these to prioritize traffic above normal, such as VoIP or important business traffic. The bottom three priorities *Limited, Limited More,* and *Limited Severely* are always limited regardless of other priorities' bandwidth consumption, because their download and upload limits are set to less than 100%. These should be used in situation where the goal is to restrict traffic regardless of if there is more bandwidth available.

Examples

Below are a few examples going from simple to more complex:

1. The network is completely idle except for one Medium priority download. This download is given all the available bandwidth and happens at full speed because no other priorities are using any traffic and the Medium download limit is 100%.
2. The network is completely idle except for one Low priority download. This download is given all the available bandwidth and happens at full speed because no other priorities are using any traffic and the Low download limit is 100%.
3. The network is completely idle except for one Limited More priority download. This download is given only half the available bandwidth because the reservation of Limited More is only 50%. The other 50% remains unused.
4. The network is fully saturated and all seven priorities have several active downloads running. All Very High Priority downloads equally split 60% of the download bandwith (the Very High reservation). All the other priorities split their reservations in a similar fashion all the way down to Limited Severely which splits the 1% reservation between all Limited Severely Sessions.
5. One Medium priority download and one Low priority download are running simultaneously. Because the other priorities are not using any of the reservations the left over is split relative to the Medium and Low reservations (5:12 or roughly 1:2.5). As such the Medium priority download runs roughly 2.5 times faster than the Low priority download and together they consume all available bandwidth. (example: Low priority runs at 100kB/sec while the Medium runs at 250kB/sec and the total available bandwidth is 350kB/sec)
6. Two Limited Severely downloads are taking place simultaneously. All sessions in the same priority share the resources of that priority so the two sessions split the priority's resources. Because all other priorities are not in use the two split the bandwidth limit (10%) and each download runs at 5% of the total available bandwidth.

7. There are two WAN interfaces and WAN Balancer is balancing traffic across the WANs. One Medium priority download is happening on WAN1 and one Low priority download is happening on WAN2. The Medium priority download is given 100% of the WAN1's bandwidth and the Low priority download is given 100% of WAN2's bandwidth.

8. There are two WAN interfaces and WAN Balancer is balancing traffic across the WANs. One Medium priority download is happening on WAN1 and one Limited More priority download is happening on WAN2. The Medium priority download is given 100% of the WAN1's bandwidth and the Limited More priority download is given 50% of WAN2's bandwidth.

> **Attention:**
>
> 1. Any given TCP download uses upload bandwidth to communicate to the sender that the data is being received. Usually this upload bandwidth is only a little, but sometimes if there is very little upload bandwidth available it can actually be the limiting factor in the total rate of the download. The receiver can only communicate with the sender to tell it that data is being received sporadically and as such the sender will slow down. This is especially common with asymmetric links, especially if other uploads are in progress.
>
> 2. Because packets are often approximately *1500* bytes (the MTU size), the lower priorities must either send the packet or not send the packet. Splitting the packet and sending a portion is not an option. As such the packet will be sent to prevent starvation but may actually exceed the 1% reservation at times. This is especially true on small links with very little bandwidth. As such the *granularity* of the limits and reservations on small links may be slightly skewed.
>
> 3. WANs are treated completely separately. Rules run on all traffic, regardless of which WAN the traffic is going out. However the bandwidth setting on each WAN is separate and they are treated as separate resources that are divided amongst traffic independently.

Settings

This section reviews the different settings and configuration options available for QoS.

- **Queue Discipline** is the queueing discipline. The queue discipline is the algorithm used to "queue" packets that are ready to be transmitted. The default is *Fair/Flow Queueing + Codel* or *fq_codel* because it is the most modern and most performant and minimizes Buffer Bloat. Another commonly used and good algorithm is *sfq* or *Stochastic Fairness Queueing* (the old default). *pfifo* or *First in First out* is the simplest but is not optimal as all packets are treated equally.

- **Enabled**: Controls whether QoS is enabled or disabled. The default setting is unchecked, which means QoS is disabled and no rules have any effect. **WAN Bandwidth should be set before enabling QoS**.

- **Default Priority**: This is the priority assigned to traffic which matches no QoS rule. It is advised to leave this at the default setting of **Medium**.

- **WAN Bandwidth**: The **most critical** setting to configure correctly. This should be set to 85-95% of your **actual line speed** - we recommend contacting your ISP to get the proper numbers and testing to verify.

  Finding the right settings for the WAN Bandwidth may take some experimentation. If the bandwidth limit is set too high QoS will have no effect at all. If the bandwidth setting is too low, traffic will be unnecessarily limited to a lower bandwidth. Remember, traffic only receives preferential treatment when the set bandwidth limit is saturated.

  The QoS limits are configured correctly when the network has slightly less throughput than when QoS is disable entirely. Depending on your hardware, the QoS settings may not match the ISP-provided or testing Mbit numbers exactly. Experimentation is required.

Figure 23: QoS - WAN Bandwidth

- **QoS Rules**: Built-in rules to provide priorities for some typically important packets and traffic types. If you're not sure, just leave these at the default.

Figure 24: QoS Rules

- **QoS Custom Rules**: This provides a simple way to create custom rules to prioritize or de-prioritize certain traffic.

  As the note warns, QoS Custom Rules **only match on bypassed traffic** - they will do **nothing** if the traffic is not bypassed. If you want to prioritize scanned sessions, use Bandwidth Control. Default rules exists for VoIP traffic, which is also bypassed by default. Here's a list of the qualifiers you can use to build Custom QoS Rules:

Name	Legal Value	Description
Destination Address	IP Matcher	The Destination IP of the traffic.
Destination Port	Int Matcher	The Destination Port of the traffic.
Destined Local		This will match on any IP the Arista holds, including aliases.Only recommended if your WAN interface(s) are Dynamic.
Protocol	Checkboxes	The protocol that should be forwarded - check all that apply.
Source Interface	Radio Buttons	The Source Interface of the traffic - choose only one.
Source Address	IP Matcher	The Source Address of the traffic.
Source MAC Address	XX:XX:XX:XX:XX:XX	The MAC Address of the source of the traffic.

- **QoS Priorities**: This table allows customization of how each priority is treated and how they are prioritized relative to other priorities. It is recommended to keep the default values.

- **Download Limit** can be any value between 1% to 100%. It limits the maximum amount of download bandwidth available to this priority under any circumstance.

- **Upload Limit** can be any value between 1% to 100%. It limits the maximum amount of upload bandwidth available to this priority under any circumstance.
- **Download Reservation** can be any value between 1% to 100%. It guarantees the minimum amount of bandwidth available to this priority should it be needed under any circumstance.
- **Upload Reservation** can be any value between 1% to 100%. It guarantees the minimum amount of bandwidth available to this priority should it be needed under any circumstance.

Figure 25: QoS Priorities

Some amount of bandwidth is always guaranteed (by the reservation) to each priority. This is to prevent any priority from being fully starved and being disconnected from the internet because higher priorities are using all the bandwidth. When a higher class is not using its reservation, the leftover is re-assigned to the lower classes based on the ratio of their reservations. For Example, by default the *Medium* priority is limited to 100% of the download bandwidth and is guaranteed at least 12% of the download bandwidth.

- **QoS Statistics** is a status readout of recent activity. The statistics are reset at reboot and when settings are saved.

It is useful for diagnosing which rules are being matched and that the proper priorities are getting assigned. It is also useful to test the total usage of each priority. The statistics are broken up by WAN interface, for example *External - Outbound* shows the priority byte counts of all traffic going out the External WAN while *External - Inbound* shows the priority byte counts of all traffic coming in the External WAN.

Interface ▲	Priority	Data
External Inbound	1 - Very High	326.25 KB
External Inbound	2 - High	1.16 GB
External Inbound	3 - Medium	37.08 MB
External Inbound	4 - Low	725 B
External Inbound	5 - Limited	0 B
External Inbound	6 - Limited More	0 B
External Inbound	7 - Limited Severely	0 B
External Outbound	1 - Very High	351.67 KB
External Outbound	2 - High	26.80 MB
External Outbound	3 - Medium	10.87 MB
External Outbound	4 - Low	503 B
External Outbound	5 - Limited	0 B
External Outbound	6 - Limited More	0 B
External Outbound	7 - Limited Severely	0 B

Figure 26: QoS Statistics

- **QoS Current Sessions** shows a table of all current active sessions and the assigned priority of each.

This is useful for testing to assure that priorities are being given the correct priority. Please note that sessions are assigned priorities at creation time, so if rules are changed active sessions will keep their current priority - only new sessions will be run against the new rules.

Related Topics

Here are several other relevant QoS links which are worthwhile reading that may give the curious administrator more insight.

- HTB (Hierarchical Token Bucket) user manual
- SFQ Stochastic Fairness Queueing

4.11.3 Access Rules

Access Filter rules apply to sessions destined to the Arista server's local processes and only sessions destined to the Arista server's local processes. These rules have no effect on sessions passing through Arista and are only used to limit and secure access to local services on the Arista server.

Note: Improperly configuring access rules can compromise the security and proper functioning of the Arista server.

Note: Disabling rules in the default configuration may interfere with the proper functioning of many features of the Arista server.

There are two rules not enabled by default:

- Allow HTTPS on WANs - enable this rule if you would like HTTPS access externally.
- Allow SSH - enable this rule if you would like SSH access to Arista's SSH service.

Note: Changing other settings of Access Rules is not recommended.

Access Rules configuration

- Enable Access Rule

  - If checked, the rule is enabled. If unchecked the rule has no effect and is disabled.
- IPv6

  - If checked, the filter rule will also be active with IPv6 addressing.
- Description

  - A description of this rule. This is just for documentation.
- Conditions

  - The conditions describing which sessions will match.
- Action

  - Block or Pass. Block means the session dropped silently. Pass means the session will be passed.

The rules are evaluated in order on all new sessions going to the Arista server as described in the Rules documentation. The action from the first matching rule is taken, if no rule matches the session is passed.

4.11.4   Universal Plug and Play

About UPnP

Universal Plug and Play (**UPnP**) allows clients to create their own firewall port forward rules. Common uses include:

- Allowing gaming consoles to host games.
- Enables BitTorrent clients to host uploads.

UPnP settings can be found at **Config** > **Network** > **Advanced** > **UPnP**.

Security Considerations

These are considered "automatic port forward rules" and therefore you should consider the potential security implications before enabling in your environment. In a home environment with an Xbox, you probably want it enabled. In an office environment, you likely do not want it enabled at all.

Settings

This section reviews the different settings and configuration options available for QoS.

- **Enabled**: Controls whether UPnP is enabled or disabled. The default setting is unchecked, which means UPnP is disabled.

- **Secure Mode**: An option that restricts port creations to the client system. In most environments you should leave this enabled.
- **UPnP Status** is a status readout of recent activity. The statistics are reset at reboot and when settings are saved.

Protocol	Client IP Address	Client Port	Destination Port	Bytes	Delete
udp	192.168.252.51	20001	20001	0.00 kbps (0.00 Mbit)	✕

C Refresh

Access Control Rules

These rules allow you to control which networks can use UPnP as well as the ports they can manage along with an allow or deny action. All rules are processed in order.

The default rules for Allow all and Deny all allow all UPnP traffic if UPnP is enabled.

If you which to control access to a particular network, create a new Allow rule for that network and ports and make sure it is above the Deny all rule.

4.11.5   Network Cards

This grid configures the various options of the network cards in the Arista server.

MTU is the maximum transmission unit. Usually the correct configuration is Auto, which in most cases is *1500*. Occasionally if the MTU is lower and PMTU (Path MTU discovery) is broken, it may be necessary to lower the MTU manually to a lower value. No settings (blank) is "Auto" - any other setting is an explicit hardcoded MTU.

The *Ethernet Media* dropdown configures the duplex of the Network card. Usually *Auto* is the correct value, however some network cards do not properly automatically configured the duplex setting or perform poorly when Auto is configured. In those cases you can manually configure the speed and duplex setting.

4.11.6   DNS and DHCP

Custom Dnsmasq Options

This textfield holds any custom Dnsmasq options. **This is for advanced users. Misconfiguration of this field will result in improper functioning of the Arista server.** Any text in this field will be appended to the dnsmasq.conf.

Dnsmasq is the server Arista uses to provide DNS & DHCP services. Dnsmasq documentation describes the options available.

Arista does not support any custom configurations for Dnsmasq.

4.11.7   Netflow

NetFlow is a feature developed by *Cisco* which provides the ability to collect IP network traffic information as it enters or exits an interface. By analyzing the data provided by NetFlow, a network administrator can determine the source and destination of traffic, class of service, and the causes of congestion. A typical flow monitoring setup consists of three main components:

- **Flow exporter**: aggregates packets into flows and exports flow records towards one or more flow collectors. In this case, the NGFW.
- **Flow collector**: responsible for reception, storage and pre-processing of flow data received from a flow exporter.
- **Analysis application**: analyzes received flow data in the context of intrusion detection or traffic profiling, for example.

Netflow on NGFW uses softflowd.

Netflow

Netflow settings are located in **Config** > **Network** > **Advanced** > **Netflow**.

- **Netflow enabled**

  - This enables the sending of netflow data to the specified netflow collector.
- **Host**

  - The IP address or hostname of the netflow collector.
- **Port**

  - The port for the netflow collector.
- **Version**

  - The version of netflow to send. NGFW supports multiple standard versions: *v1*, *v5*, and *v9*.

4.11.8   Dynamic Routing

About Dynamic Routing

Dynamic Routing allows for the exchange of routes between other routers using Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF).

Dynamic Routing settings can be found at **Config** > **Network** > **Advanced** > **Dynamic Routing**.

BGP Overview

The Border Gateway Protocol (BGP) requires that all nodes known as *neighbors* are known and added to the settings.

OSPF Overview

Open Shortest Path First (OSPF) does not require all nodes to be known. Instead each route is associated with a group called an *area*. OSPF can be hierarchical through the use of multiple areas so some networks are publicly known and others can be private. Additionally OSPF supports authentication.

Settings

This section reviews the different settings and configuration options available for Dynamic Routing.

- **Dynamic Routing Enabled**: Controls whether dynamic routing is enabled or disabled. The default setting is unchecked, which means dynamic routing is disabled. BGP and/or OSPF must also be enabled as well.
- **Status**: Overall status of dynamic routing shows:

  - **Acquired Dynamic Routes**: All routes obtained from enabled dynamic routing protocols.
  - **BGP Status** Information about each BGP neighbor including messages received, sent, and uptime.

- **OSPF Status** Information about discovered OSPF neighbor such as their IP address and the time remaining until they next synchronize.

- **BGP**: Enable BGP protocol.

  - **Router ID**: An IP-like identifier. It can be any IP-address like value but is typically your WAN address.
  - **Router AS**: The Autonomous System (AS) number for this system. It can be any number from *1-65535* but must be unique in your BGP network.

  - **Neighbors** Define each BGP neighbor here. For each neighbor, you will need to know their IP address and AS.
  - **Networks** Define each local network route to share via BGP.
- **OSPF**: Enable OSPF protocol.

4.12   Network Reports

Network reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports

Reports can be searched and further defineNetwork Reportsd using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
Network Summary	A summary of network traffic.
Data Usage (by interface)	The total data usage by interface.
Data Usage per Day (by interface)	The data usage of each interface by day
Data Rx-Usage (by interface)	The total received data usage by interface.
Data Tx-Usage (by interface)	The total received data usage by interface.
Sessions	The amount of total, scanned, and bypassed sessions over time.
Sessions Per Minute	The amount of total, scanned, and bypassed sessions created per minute.
Sessions Per Hour	The amount of total, scanned, and bypassed sessions created per hour.
Bandwidth Usage	The approximate averaged data transfer rate (total, sent, received) over time.
Top Client Addresses	The number of sessions grouped by client (source) address.
Top Server Addresses	The number of sessions grouped by server (destination) address.
Top Server Ports	The number of sessions grouped by server (destination) port.
Top IP Protocols	The number of sessions grouped by IP protocol number.
Top Server Countries	The number of sessions grouped by server (destination) country.
Interface Usage	The RX rate of each interface over time.
All Sessions	All sessions handled by Arista.
Scanned Sessions	All sessions that were not bypassed.
Bypassed Sessions	All sessions matching a bypass rule and bypassed.
Blocked Sessions	All sessions blocked by filter rules.
Port Forwarded Sessions	All sessions match a port forward rule.
NATd Sessions	All sessions that have been NATd by Arista.
All Session Minutes	All sessions by minute.

The tables queried to render these reports:

- Sessions from the Database Schema

4.13   Troubleshooting

Setting	Description
Connectivity Test	The Connectivity Test checks that your NG Firewall can resolve and connect to http://updates.arista.com. This is an important test to establish that your WAN connections is functioning properly.
Ping Test	A simple Ping utility. Enter a hostname or IP and ping away.
DNS Test	A simple DNS utility. Enter a hostname and get an IP.
Connection Test	The Connection Test is a very useful tool that lets you check the status of a port on a remote machine.Enter an IP or Hostname and a Port, click **Run Test**, and see what happens.
Traceroute Test	A Simple Traceroute utility. Enter a hostname or IP and see what's between your NG Firewall and the remote machine.
Download Test	Provides a utility to test the download speed that NG Firewall has available. This download is not scanned so this will provide an upper bound of the bandwidth available to NG Firewall itself as a single TCP download.Please note that the results of this test are displayed in MBps, megabytes per second. You will need to multiply this result value by eight to get the speed test result in Mbps, megabits per second.
Packet Test	The Packet Test is a very powerful troubleshooting tool. Select an Interface to listen on and a timeout value, then hit Run Test to see traffic on that interface.You can filter by IP and/or port to, for example, check if traffic is hitting an interface or if a remote machine is answering a request.

NG Firewall Performance Apps

Contents

5.1 Bandwidth Control

Bandwidth Control gives you the power to monitor and control bandwidth usage on your network.

About Bandwidth Control

Bandwidth control can be used to ensure that your network continues to operate smoothly and that bandwidth is shared optimally based on what is important to you. Many organizations struggle with bandwidth problems such as students watching online videos or clients using BitTorrent while more important tasks struggle to complete for bandwidth. You can use Bandwidth Control to do things like give high priority to Video Chat or slow down all traffic coming out of machines using BitTorrent.

> **Note:** Enabling Bandwidth Control automatically enables QoS. But disabling Bandwidth does not automatically disable QoS.

Settings

This section reviews the different settings and configuration options available for Bandwidth Control.

Status

This displays the current status and some statistics

Setup Wizard

The setup wizard configures the initial configuration of Bandwidth Control - please pay attention to the prompts as they provide valuable information on how the application works and the answers to your questions will determine the configuration.

- **Configure WAN download and upload bandwidth:** After the welcome screen, you will be asked to set the bandwidth rates of your WAN interface. This is the most important setting in the configuration of Bandwidth Control. If you are unsure it is recommended to run some bandwidth tests when there is no other activity to determine the true download and upload rates of your WAN connection. Entering a value around 95%-100% of the measured value is typically ideal. If the value is too low, Bandwidth Control will unnecessarily limit bandwidth to the value you have entered. If the value is too high, Bandwidth Control will be less effective as it will over-allocate bandwidth and lose some ability to differentiate by priority. You will be asked to repeat this process for each WAN interface.
- **Choose a starting configuration:** After setting the WAN settings, choose a starting configuration that best suites your organization. Each configuration's goals are described as well as what is prioritized and deprioritized. These rules can be customized later - this is just a starting configuration.
- **Quotas:** In addition to the starting configuration, quotas can also be configured. Most sites will not need quotas, however quotas can be extremely useful in some scenarios to prevent users from monopolizing resources. To enable quotas, click on Enable and provide information that best suites your organization.
    - **Quota Clients:** The clients will be given quotas. Be careful to not give a range that includes any servers and machines that you don't want to have quotas.
    - **Quota Expiration:** The expiration time of each quota (or length of time the quota will be in use.) After a quota expires a new quota will be granted.
    - **Quota Size:** The size of the quota each host is granted (in bytes).
    - **Quota Exceeded Priority:** The priority given to hosts after they exceed their quota (if they do so).

More information on Quotas and how they work can be found in the Quotas section.

After this your configuration of Bandwidth Control is complete and Bandwidth Control is enabled!

Rules

The rules tab contains most of the configuration and settings controlling the behavior of Bandwidth Control. Rules determine the action that will be taken when traffic passes through Bandwidth Control.

For each session the rules are evaluated in order until the first match is found, then the action associated with the matching rule is performed and the data chunk is sent on its way. If no rule is found the no action is taken. If the session has been given no priority it is given the default QoS priority, which is normally Medium.

> **Note:** Unlike most Rules in other apps, the rules in Bandwidth Control are consulted not only when the session is formed but also again on the first ten packets because some matchers such as "HTTP: Hostname" or "Application Control: Application" are not known until several packets into the session. Also, all of a host's sessions will be reevaluated when they are added/removed to the penalty box or when a quota is exceeded so active sessions will be reprioritized accordingly.

Extensive rule sets can be created (and imported and exported) that carefully assign the correct priorities to the desired traffic and perform the desired actions at the desired times.

The Rules documentation describes how rules work and how they are configured.

Rule Actions

- **Set Priority** Sets the matching session to the chosen priority.

  - *Priority* The priority to be assigned.
- **Tag Host** adds tag to the host to mark it for further actions.
- **Give Host a Quota:** Gives the host IP a quota.

  - **Quota Expiration** defines how long their quota will last

    - "End of Hour" means the quota will expire at the 59th minute of the hour.
    - "End of Day" means the quota will expire at 11:59pm of the day.
    - "End of Week" means the quota will expire 1 minute before the end of week (Saturday 11:59pm if US-localized)
    - An integer can also be specified for the number of seconds the quota will last from the creation date.
  - **Quota Bytes** defines the number of bytes in their quota.
- **Give User a Quota:** Gives the user a quota

  - **Quota Expiration** defines how long their quota will last

    - "End of Hour" means the quota will expire at the 59th minute of the hour.
    - "End of Day" means the quota will expire at 11:59pm of the day.
    - "End of Week" means the quota will expire 1 minute before the end of week (Saturday 11:59pm if US-localized)
    - An integer can also be specified for the number of seconds the quota will last from the creation date.
  - **Quota Bytes** defines the number of bytes in their quota.

	Rule Id	Enable	Description	Conditions	Action	Edit	Delete
	100001	☑	Give User a Quota if no Quota	User has no Quota ⇒ **True** · Source Interface ⇒ **Any Non-WAN**	Give User a Quota		🗑
	100002	☑	Penalize Users over Quota	User has exceeded Quota ⇒ **True**	Set Priority [Limited More]		🗑
	100003	☑	Give Host a Quota if no Quota	Host has no Quota ⇒ **True** · Source Interface ⇒ **Any Non-WAN**	Give Host a Quota		🗑
	100004	☑	Penalize Hosts over Quota	Host has exceeded Quota ⇒ **True**	Set Priority [Limited More]		🗑
	100005	☑	Apply Penalty Box Penalties	Tagged ⇒ **penalty-box**	Set Priority [Limited Severely]		🗑
	100006	☑	Prioritize DNS	Destination Port ⇒ **53**	Set Priority [Very High]		🗑
	100007	☑	Prioritize SSH	Protocol ⇒ **TCP** · Destination Port ⇒ **22**	Set Priority [High]		🗑
	100008	☑	Prioritize Remote Desktop (RDP,VNC)	Protocol ⇒ **TCP** · Destination Port ⇒ **3389, 5300**	Set Priority [Very High]		🗑
	100009	☑	Prioritize eMail (POP3,POP3S,IMAP,I…	Destination Port ⇒ **110, 995, 143, 993**	Set Priority [High]		🗑
	100010	☑	Prioritize "Remote Access" traffic (r…	Application Control: Application Category ⇒ **Remote Access**	Set Priority [High]		🗑
	100011	☑	Deprioritize "Unproductive" Applicat…	Application Control: Productivity ⇒ **<2**	Set Priority [Low]		🗑
	100012	☑	Deprioritize site violations (require…	Web Filter: Website is Flagged ⇒ **True**	Set Priority [Medium]		🗑
	100013	☑	Deprioritize Windows updates (dow…	HTTP: Hostname ⇒ ***windowsupdate.com**	Set Priority [Low]		🗑
	100014	☑	Deprioritize Microsoft updates (upd…	HTTP: Hostname ⇒ ***update.microsoft.com**	Set Priority [Low]		🗑
	100015	☑	Deprioritize dropbox.net sync	HTTP: Hostname ⇒ ***dropbox.com**	Set Priority [Low]		🗑
	100016	☑	Tag Bittorrent users for 30 minute…	Application Control: Application ⇒ **BITTORRE**	Tag Host		🗑
	100017	☑	Deprioritize P2P traffic (requires Ap…	Application Control Lite: Category ⇒ **Peer to Peer**	Set Priority [Low]		🗑
	100018	☑	Deprioritize File Transfers (requires…	Application Control: Application Category ⇒ **File Transfer**	Set Priority [Low]		🗑
	100019	☑	Deprioritize HTTP to Download Site…	Web Filter: Category ⇒ **Download Sites**	Set Priority [Low]		🗑
	100020	☑	Limit dropbox.com sync	HTTP: Hostname ⇒ ***dropbox.com**	Set Priority [Limited More]		🗑
	100021	☑	Do not Prioritize large HTTP downlo…	HTTP: Content Length ⇒ **>10000000**	Set Priority [Medium]		🗑
	100022	☑	Prioritize HTTP	Destination Port ⇒ **80**	Set Priority [High]		🗑

Priorities

The overall effect of Bandwidth Control is to map traffic to priorities which are enforced by the QoS engine. There are 7 Priorities: Very High, High, Medium, Low, Limited, Limited More, and Limited Severely.

The first four priorities can be thought of as "normal" - Very High, High, Medium, and Low. They are given certain precedence over bandwidth rights. Very High traffic has the option to consume bandwidth before High, Medium, and Low. The Very High bucket will be assigned the largest amount of bandwidth, less to High, even less to Medium, and much less to Low.

The other three - Limited, Limited More, and Limited Severely - are different in that they will never use all available bandwidth. The classes are punitive because they will limit bandwidth to a percentage of the whole even if there is more available.

To read much more in depth about the effects of prioritization and how bandwidth allotment works, see QoS.

Note: Effective Bandwidth Shaping is all about assigning the correct priorities such that important traffic is never starved by less important traffic.

A fundamental principle is that limiting traffic to a fixed low rate enforcement is almost never the right thing to do because wasted bandwidth is irrecoverable. In cases where the desire is to starve less important traffic it should be assigned a lesser priority (medium or low) so that it can still consume all bandwidth if no more important tasks are available. This means the less important task will finish quicker so that later these resources are free and this occurs definitionally at no expense to higher priority traffic, ever.

The priorites that limit to less than 100% even when the bandwidth is unused (Limited, Limited More, and Limited Severely by default) are useful for punitive situations.

Quotas

Quotas are set amounts of data that can be used over a certain amount of time. This is useful for sites where you want to punish excessive usage. For example, in a hotel we want each IP to get 1 GB a day, but if this amount is exceeded it will be considered excessive and that host can be treated differently (be blocked, receive less bandwidth, etc). By using quotas and rules, bandwidth abusers are handled automatically requiring no administrator intervention.

Quotas can be assigned to Users or Hosts and the current quota status can be viewed by clicking on Users or Hosts accordingly. All sessions' data passing through NG Firewall gets counted against the corresponding Host or User.

Reports

The Reports tab provides a view of all reports and events for all traffic handled by Bandwidth Control.

Reports

This applications reports can be accessed via the Reports tab at the top or the Reports tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the Conditions window at the bottom of the page. The data used in the report can be obtained on the Current Data window on the right.

Pre-defined report queries:

Report Entry	Description
Bandwidth Control Summary	A summary of Bandwidth Control actions.
Bandwidth Usage	The approximate averaged data transfer rate (total, sent, received) over time.
Top Hostnames Usage	The bandwidth usage of the top hostnames.
Top Hostnames (by total bytes)	The sum of the data transferred grouped by hostname.
Top Hostnames (by received bytes)	The sum of the received data grouped by hostname.
Top Hostnames (by sent bytes)	The sum of the sent data grouped by hostname.
Top Clients Usage	The bandwidth usage of the top clients.
Top Clients (by total bytes)	The sum of the data transferred grouped by client address.
Top Usernames Usage	The bandwidth usage of the top usernames.
Top Usernames (by total bytes)	The sum of the data transferred grouped by username.
Top Server Port Usage	The bandwidth usage by top server port.
Top Ports (by total bytes)	The sum of the data transferred grouped by server port.
Top Ports (by received bytes)	The sum of the data received grouped by server port.
Top Ports (by sent bytes)	The sum of the data sent grouped by server port.
Top Applications Usage	The bandwidth usage of the top applications.
Top Application (by total bytes)	The sum of the data transferred grouped by Application Control application.
Top Application (by received bytes)	The sum of the data sent grouped by Application Control application.
Top Application (by sent bytes)	The sum of the data sent grouped by Application Control application.
Top Categories Usage	The bandwidth usage of the top application categories.
Top Category (by total bytes)	The sum of the data transferred grouped by Application Control category.
Top Priorities Usage	The bandwidth usage by priority.

Report Entry	Description
Top Priorities (by total bytes)	The sum of the data transferred grouped by priority.
Top Countries Usage	The bandwidth usage by top countries.
Top Countries (by total bytes)	The sum of the data transferred grouped by country.
Bypassed (by total bytes)	The sum of the data transferred grouped by bypassed.
All Sessions	All sessions processed by Bandwidth Control.
Quota Events	Shows when quotas are assigned or expired.
Prioritized Sessions	All sessions prioritized by Bandwidth Control.

The tables queried to render these reports:

- Sessions
- Database Schema

Related Topics

- QoS

5.1.1 Bandwidth Control Reports

The Reports tab provides a view of all reports and events for all traffic handled by Bandwidth Control.

Reports

This applications reports can be accessed via the Reports tab at the top or the Reports tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created. Reports can be searched and further defined using the time selectors and the Conditions window at the bottom of the page. The data used in the report can be obtained on the Current Data window on the right. Pre-defined report queries:

Report Entry	Description
Bandwidth Control Summary	A summary of Bandwidth Control actions.
Bandwidth Usage	The approximate averaged data transfer rate (total, sent, received) over time.
Top Hostnames Usage	The bandwidth usage of the top hostnames.
Top Hostnames (by total bytes)	The sum of the data transferred grouped by hostname.
Top Hostnames (by received bytes)	The sum of the received data grouped by hostname.
Top Hostnames (by sent bytes)	The sum of the sent data grouped by hostname.
Top Clients Usage	The bandwidth usage of the top clients.
Top Clients (by total bytes)	The sum of the data transferred grouped by client address.
Top Usernames Usage	The bandwidth usage of the top usernames.
Top Usernames (by total bytes)	The sum of the data transferred grouped by username.
Top Server Port Usage	The bandwidth usage by top server port.

Report Entry	Description
Top Ports (by total bytes)	The sum of the data transferred grouped by server port.
Top Ports (by received bytes)	The sum of the data received grouped by server port.
Top Ports (by sent bytes)	The sum of the data sent grouped by server port.
Top Applications Usage	The bandwidth usage of the top applications.
Top Application (by total bytes)	The sum of the data transferred grouped by Application Control application.
Top Application (by received bytes)	The sum of the data sent grouped by Application Control application.
Top Application (by sent bytes)	The sum of the data sent grouped by Application Control application.
Top Categories Usage	The bandwidth usage of the top application categories.
Top Category (by total bytes)	The sum of the data transferred grouped by Application Control category.
Top Priorities Usage	The bandwidth usage by priority.
Top Priorities (by total bytes)	The sum of the data transferred grouped by priority.
Top Countries Usage	The bandwidth usage by top countries.
Top Countries (by total bytes)	The sum of the data transferred grouped by country.
Bypassed (by total bytes)	The sum of the data transferred grouped by bypassed.
All Sessions	All sessions processed by Bandwidth Control.
Quota Events	Shows when quotas are assigned or expired.
Prioritized Sessions	All sessions prioritized by Bandwidth Control.

The tables queried to render these reports:

- Sessions
- Session Minutes

Related Topics

Report Viewer

Reports

5.2   Branding Manager

Branding Manager is designed to allow you to rebrand user-facing components by adding your company logo, name, URL and contact email.

Branding Manager will replace the "Arista Edge Threat Management" branding in all user-facing interactions, such as block pages, quarantine digest emails, quarantine digests, root certificate installer, etc. This is not meant to remove all Arista Edge Threat Management branding; the administrator UI still contains many references to NG Firewall. For HTML/CSS experts, combining Branding Manager with a custom block page skin and a custom rack skin gets you a fully customized style. Of course, you don't need a custom skin to change any of the branding elements - see the table below to see what you can change with Branding Manager!

Settings

Status

This displays the current status and some statistics.
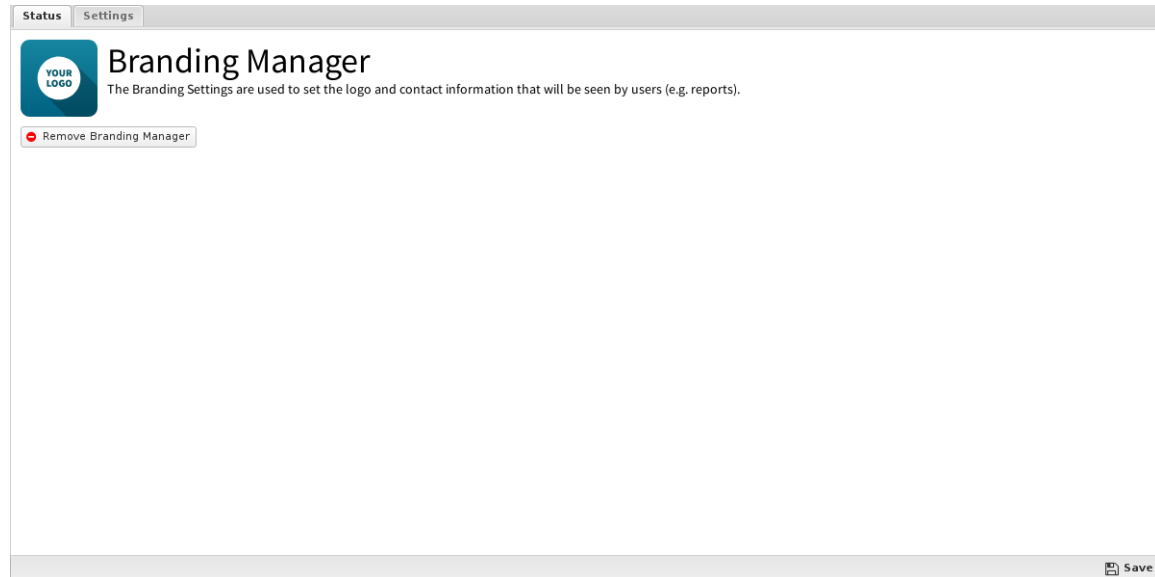
Settings

This section reviews the different settings and configuration options available for Branding Manager.

To modify your Branding Manager settings, hit the **Settings** button on the faceplate. This will bring up a menu where you upload a custom logo as well as modify name, url and contact settings. Note that the recommended resolution for logos is **150x100**; the maximum resolution is **166x100**. All image formats are supported, however we do not recommend including animation as it can affect the PDF reports. Branding Manager's text fields have a 256 character limit.

- **Logo:** Use this option to upload a replacement logo.
- **Contact Information**
  - **Company Name:** The name of your company.
  - **Company URL:** The URL for your website (for example http://www.arista.com).
  - **Contact Name:** The name of the network administrator that should be contacted if questions or problems arise.
  - **Contact Email:** The email address of the network administrator (for example, user@domain.com)
- **Banner Message**
  - **Message Text:** Text that will be displayed above login boxes. This is restricted to plain text.

Fine-tune Your Logo

It is not required that you fine-tune your logo, but it will make it look more professional. Here are a few tips:

- Convert your image to Greyscale: Using Adobe Photoshop, open your image and go to **File** > **Save As Web & Devices**, specify **GIF** and **Greyscale**, then save.
- Resize your image: Using Adobe Photoshop, open your image and go to **File** > **Save As Web & Devices**, click the **Image** tab and resize to 150x100.
- Place your logo on a background: Download the templates, then overlay your logo.

5.3   WAN Balancer

WAN Balancer works in conjunction with multiple ISPs to distribute your traffic across multiple connections. It will decide dynamically which WAN connection to send traffic over, maximizing your bandwidth usage.

You may also consider using WAN Failover in your network as well - it automatically reroutes traffic over working WAN links when one fails. If WAN Failover is running and detects a WAN as being down WAN Balancer will not balance traffic to that WAN.

Settings

This section reviews the different settings and configuration options available for WAN Balancer.

Status

This tab displays information and statistics for each WAN interface.

Traffic Allocation

On the **Traffic Allocation** tab, you set the weighs for each WAN connection. If you only have one WAN defined, you will only see one interface listed here. Simply enter the weighting you desire, check that you are good with the percentages assigned to each WAN, and hit save.

As each WAN Balancer processes each new session it decides which WAN it will use to send this traffic if there is no local route for the traffic. If traffic between these two IPs has taken place recently then there is likely a route already in the cache. If so, this route will be used to send this new session. This is to assure that all traffic between two IPs uses the same WAN consistently to avoid issues with cloud services.

If there is no route in the cache, then a WAN will be chosen based on the hash of the source and destination and the weights given in the *Traffic Allocation* settings. Not that the traffic allocation weights don't determine exactly the percentages of traffic over the various WANs, only how sessions will be assigned to various WANs.

Status | **Traffic Allocation** | Route Rules | 📊 View Reports

Allocate traffic across WAN interfaces

Traffic allocation across WAN interfaces is controlled by assigning a relative weight (1-100) to each interface.
After entering the weight of each interface the resulting allocation is displayed.
If all WAN interfaces have the same bandwidth it is best to assign the same weight to all WAN interfaces.
If the WAN interfaces vary in bandwidth, enter numbers that correlate the relative available bandwidth.
For example: 15 for a 1.5Mbit/sec T1, 60 for a 6 mbit link, and 100 for a 10mbit link.

Interface Weights

Interface	Weight	Resulting Traffic Allocation
External	50	100% of Internet traffic.

💾 Save

Route Rules

Route Rules determine which WAN will be used for traffic going to the internet (traffic with no local route). As described in the Rules documentation, the *Route Rules* are evaluated in order for new sessions and the first matching rule will determine which WAN interface is used. If no matching rule is found or the first matching rule has a *Destination WAN* set to *Balance* then the session will be randomly assigned a route based on *Traffic Allocation* settings. A limited set of conditions are available to WAN Balancer Route Rules which include source, destination, port and protocol.

This allows you to specify which WAN is used for certain traffic based on a variety of conditions. For example:

- To put all one server's traffic on a specific WAN add a rule with condition "*Source Address* is server_ip" and the *Destination WAN* is the WAN to be used.
- To put all SMTP to go out a specific WAN, add a rule with "*Destination Port* is **25**" and the *Destination WAN* is the WAN to be used.

This can also be useful if you have one connection with less throughput but lower latency. In this case you can specify that all VOIP or latency sensitive traffic use the lower latency connection.

✏️ **Note:** Unlike Routes, Route Rules that route traffic to a WAN that is down will automatically balance traffic to one of the active WANs. For example, if a rule says to send all **port 25** traffic to WAN2, but WAN Failover knows WAN2 is down, this rule will effectively mean *Balance* which means the session will be put on one of the other active WANs.

✏️ **Note:** Routes and routes based on the network configuration always override Route Rules. Route Rules only apply to sessions that have no local route based on configuration on routes in Routes. Route Rules are basically a suggestion that the traffic be routed out a specific WAN if no other route says where to send it.

Reports

The Reports tab provides a view of all reports and events for all traffic handled by WAN Balancer.

This applications reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
WAN Balancer Summary	A summary of WAN Balancer actions.
Sessions By Interface	The number of sessions destined to each interface.
Bytes By Interface	The number of bytes destined to each interface.

The tables queried to render these reports:

- Database Schema

Related Topics

- WAN Failover

5.3.1   WAN Balancer Reports

The Reports tab provides a view of all reports and events for all traffic handled by WAN Balancer.

Reports

This applications reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
WAN Balancer Summary	A summary of WAN Balancer actions.
Sessions By Interface	The number of sessions destined to each interface.
Bytes By Interface	The number of bytes destined to each interface.

The tables queried to render these reports:

- Database Schema

Related Topics

Report Viewer

Reports

5.4 WAN Failover

WAN Failover works in conjunction with multiple ISPs to assure that you maintain Internet connectivity if a loss of connectivity occurs on one of your WAN connections. If one of your ISP links goes down, WAN Failover will automatically route all traffic over the other WAN(s) until service is restored.

You may also consider using WAN Balancer in your network as well - it allows you to maintain an automatic distribution of traffic over multiple WAN links rather than just failing over if one goes down.

Tests are configured for each WAN which are run continuously to determine the current status of each interface. If enough test fail on a given WAN to exceed the failure threshold then the WAN is considered down and internet-bound traffic will not go out that WAN. The lowest ID active WAN is used as the current default WAN interface for internet-bound traffic.

Settings

This section reviews the different settings and configuration options available for WAN Failover.

Status

The Status tab presents an overview of the WANs and the current test results.

- **Interface ID**: The number of the interface.
- **Interface Name**: The name of the interface in the NG Firewall GUI.
- **System Name**: The name of the interface as seen by NG Firewall.
- **Online Status**: True or False whether the WAN is online.
- **Current Tests Count**: The total number of the tests ran on that interface.
- **Tests Passed**: The total number of the tests ran on that interface that passed.
- **Tests Failed**: The total number of the tests ran on that interface that failed.

Tests

WAN Failover must have tests set up for every WAN interface; these tests are set up on the Tests tab. Just click **Add**, select your interface and test type, then run the test - if it passes, go ahead and save it.

Tests is how WAN Failover determines if the given WAN interface is up or down so it is important to pick tests carefully that correlate with the status of that WAN connection. For example, pinging an ISP router is generally a good test because it will usually fail when the ISP is down but work when connectivity is good. Pinging a public site like google.com may work, but may sometimes have false positives or false negatives. Pinging the gateway may also work, but may sometimes provide false positives when the gateway is reachable but the ISP itself is offline.

The options are as follows:

* **Interface**: The interface you want to set up a test for.
* **Description**: A description for this test.
* **Testing Interval**: Determines how often (in seconds) your specified test will be executed.
* **Timeout**: The maximum amount of time that may pass without receiving a response to your test. This value should be less than the **Testing Interval**. You should make sure that you allow for enough time to pass if you have a poor connection to the internet, or a connection that often has long latency (delays) associated with it.
* **Failure Threshold**: How many failures are acceptable during the testing period.
* **Test Type**: is the specific method you will use to determine whether failover will be initiated.

Note on DNS tests:

> ⚠ **Warning:** DNS tests use all the DNS entries in the Interface WAN settings. If the DNS entries are only available on a specific WAN, for example ISP DNS only available on their network, then routes must be configured for those DNS servers. Otherwise, some DNS tests will fail as the DNS is not reachable on a non-ISP WAN making NG Firewall falsely see the WAN as down.
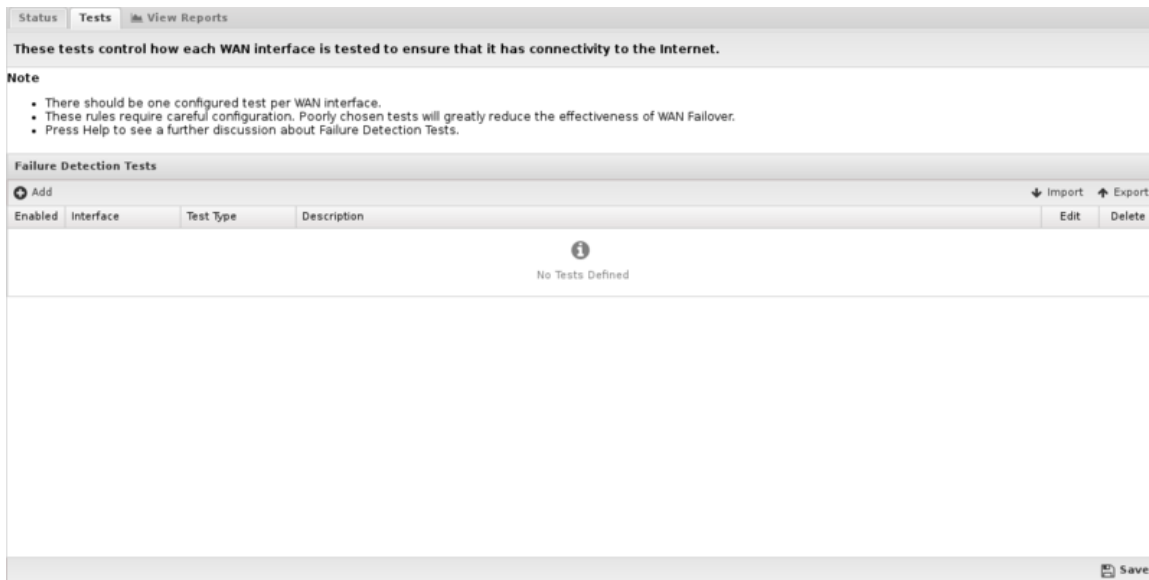
Reports

The Reports tab provides a view of all reports and events for all traffic handled by WAN Failover.

This applications reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
WAN Failover Summary	A summary of WAN Failover actions.
WAN Disconnect Events	The number of disconnect events grouped by WAN.
WAN Interface Outages	The fails tests of each interface over time.
Outage Events	Events where the failure threshold was exceeded and the WAN was considered offline.
Test Events	All test events and their outcome.
Failed Test Events	All tests that resulted in failure.
Success Test Events	All tests that resulted in success.

The tables queried to render these reports:

- WAN Failover Reports

Related Topics

- WAN Balancer

5.4.1   WAN Failover Reports

The Reports tab provides a view of all reports and events for all traffic handled by WAN Failover.

Reports

This applications reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
WAN Failover Summary	A summary of WAN Failover actions.
WAN Disconnect Events	The number of disconnect events grouped by WAN.
WAN Interface Outages	The fails tests of each interface over time.
Outage Events	Events where the failure threshold was exceeded and the WAN was considered offline.
Test Events	All test events and their outcome.
Failed Test Events	All tests that resulted in failure.
Success Test Events	All tests that resulted in success.

The tables queried to render these reports:

- Sessions
- Session Minutes

Related Topics

Report Viewer

Reports

5.5 Web Cache

Web Cache application provides HTTP content caching: as web traffic passes through the NG Firewall server it will be transparently cached. This will both save bandwidth by serving repeat content from the local cache and improve user experience by loading cached sites faster.

Just like Web Filter and other applications on NG Firewall, Web Cache works transparently on traffic passing through the NG Firewall server. There is no need to change any of the settings on any of the PCs behind NG Firewall to gain the benefits of web caching.

As content is downloaded from the web it is stored in a local cache on the disk. Upon later requests of the same web document the content is served directly from the local cache. The same document does not get downloaded multiple times, and the client gets a better user experience because they don't have to wait on subsequent downloads of the same document.

Cache Bypass

The **Cache Bypass** tab allows you to enter sites you do not want cached. Some sites to not operate properly with Web Cache working (such as Google Maps, which is bypassed by default) so you may need to add some sites to this list. Just hit Add, fill in the domain name, and save.

Reports

The Reports tab provides a view of all reports and events for all traffic handled by Web Cache.

This applications reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
Web Cache Summary	A summary of Web Cache actions.
Cache Hit-Miss Statistics	The number of cache hits, misses, and sessions bypassed over time.
Cache Size Statistics	The amount of cached and uncached web data over time.
Web Cache Events	All HTTP events processed by Web Cache.

The tables queried to render these reports:

- Web cache stats in the Database Schema

Related Topics

Web Filter

Settings

This section reviews the different settings and configuration options available for Web Cache.

Status

The **Status** tab displays statistics from Web Cache - there are no settings to configure.

- **Statistics**: The following information will help you understand the statistics Web Cache provides:

    - **Cache Hit Count** displays the total number of HTTP requests that have been served from cache.
    - **Cache Miss Count** displays the total number of HTTP requests that were not found in cache and thus were served using content retrieved from the external server where the content resides.
    - **Cache Hit Bytes** displays the size, in bytes, of all HTTP requests that have been served from cache.
    - **Cache Miss Bytes** displays the size, in bytes, of all HTTP requests that were not found in cache.
    - **User Bypass Count** displays the number of HTTP sessions that bypassed the cache because the server hosting the content was listed in the user managed cache bypass list.
    - **System Bypass Count** displays the number of HTTP sessions that bypassed the cache because the system determined they were not compatible with our caching model. Web Cache can generally handle all GET and HEAD requests, and we also allow smaller POST requests to transit through the cache logic. Everything else (ie: Large POST requests, non HTTP traffic, etc.) will be allowed to bypass the cache entirely.

- **Clear Cache**: If content stored in the cache somehow becomes stale or corrupt, the cache can be cleared with the Clear Cache button. As noted in the GUI, clearing the cache requires restarting the caching engine, which will cause active web sessions to be dropped and may disrupt web traffic for several seconds.

5.5.1   Web Cache Reports

The Reports tab provides a view of all reports and events for all traffic handled by Web Cache.

Reports

This applications reports can be accessed via the *Reports* tab at the top or the *Reports* tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the *Conditions* window at the bottom of the page. The data used in the report can be obtained on the *Current Data* window on the right.

Pre-defined report queries:

Report Entry	Description
Web Cache Summary	A summary of Web Cache actions.
Cache Hit-Miss Statistics	The number of cache hits, misses, and sessions bypassed over time.
Cache Size Statistics	The amount of cached and uncached web data over time.
Web Cache Events	All HTTP events processed by Web Cache.

The tables queried to render these reports:

- Database Schema

Related Topics

Report Viewer

Reports

NG Firewall Connect Apps

Contents

6.1   Captive Portal

Captive Portal allows administrators to require network users to log in or accept a network usage policy before accessing the internet.

Captive Portal can authenticate users against NG Firewall's built-in Local Directory, Active Directory (if Directory Connector is installed), or RADIUS. It can be used to set up policies (for Policy Manager) by username (or group name if using Active Directory) rather than IP. While Captive Portal is running, **captured** machines will be forced to authenticate (or just press OK) on the Captive Portal page before they are able to access the internet.

Captive Portal is a common technique used to identify users on the network as describe in Users.

Getting Started with Captive Portal

After installing Captive Portal, complete the following steps to get it working:

1. Define which machines will be **captured** and required to complete the Captive Portal process before accessing the Internet - enabling the first example rule in the Capture Rules table will force all machines on the internal interface to use Captive Portal.
2. Enter any IPs that unauthenticated machines will need to access - these can be entered in the **Pass Listed Server Addresses** section of the **Passed Hosts** tab.
3. Enter any IPs that always need access to the internet - these can be entered in the **Pass Listed Client Addresses** section of the **Passed Hosts** tab.
4. Customize the Captive Portal page on the **Captive Page** tab. If **Basic Login** is chosen, set the appropriate authentication method for users on the **User Authentication** tab.
5. Turn on Captive Portal.

At this point Captive Portal will evaluate your **Capture Rules** and any traffic that matches will be stopped until that user has completed the Captive Portal process.

Settings

This section reviews the different settings and configuration options available for Captive Portal.

Status

This tab shows the current status of Captive Portal. You can see information about current captured IPs, such as the username and other session information. You can also logout any active session.
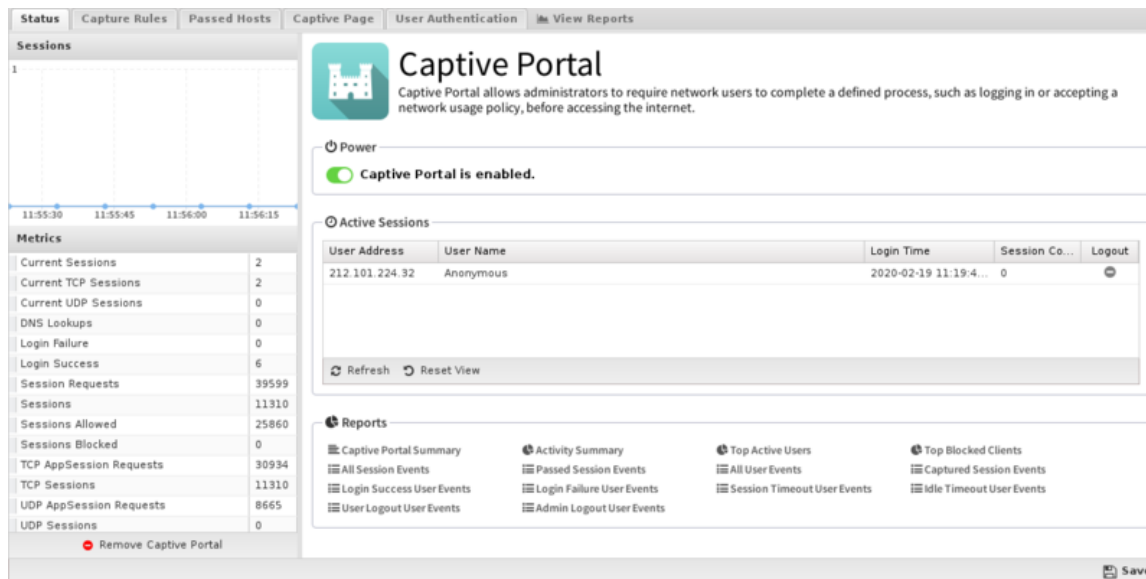
Figure 27: Captive Portal Status

Capture Rules

The **Capture Rules** tab allows you to specify rules to Capture or Pass traffic that crosses the NG Firewall.

The Rules describes how rules work and how they are configured. Captive Portal uses rules to determine whether to capture or pass each network session. The rules are evaluated in order, and on the first match, the configured action will be applied. If no rules match, the traffic is allowed by default.

If the action is *Pass* the session is passed, regardless of the clients authentication status. If the action is *Capture* the session is "captured" which means several different things depending on several factors:

- If the client is authenticated, the session is passed
- If the client is not authenticated and the protocol is tcp and the destination port = 80, then a redirect to the captive portal page is sent.
- If the client is not authenticated and the protocol is tcp and the destination port = 443, then a redirect to the captive portal page is sent. (The certificate will not match as the captive portal is not actually the requested server)
- If the client is not authenticated and the destination port = 53, then a DNS response is sent after validating it is a valid DNS request.
- If the client is not authenticated and the session has a destination port != 53,80,443 then then session is blocked.
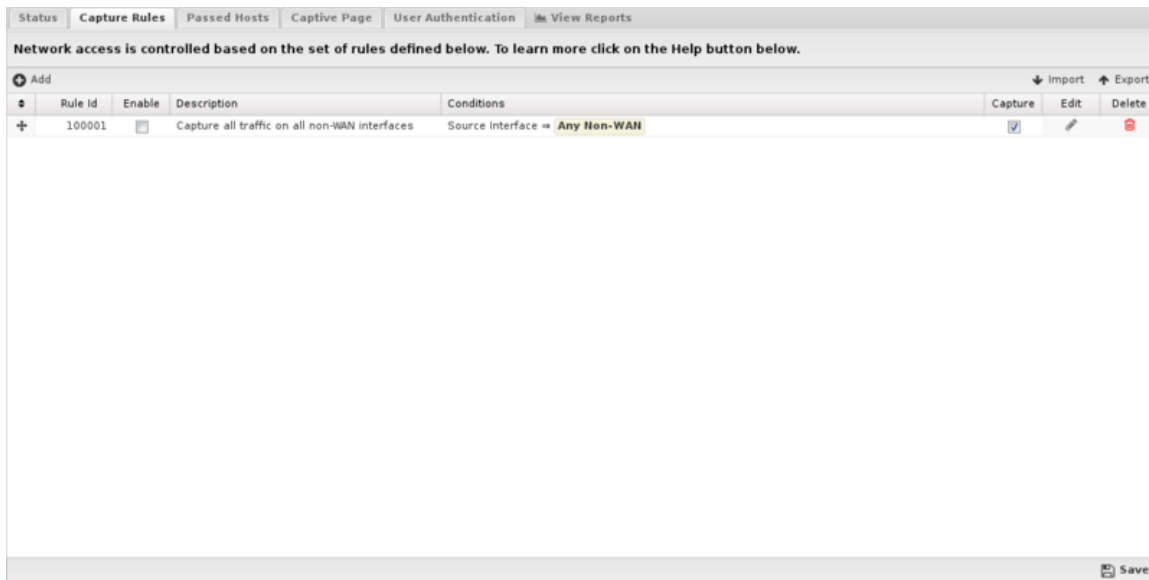
Figure 28: Capture Rules Tab

Passed Hosts

The **Pass Hosts** tab allows you to specify machines that either **a)** should not be affected by Captive Portal, or **b)** servers that machines behind Captive Portal should be able to access even if unauthenticated.

- **Pass Listed Client Addresses**: These machines will not be affected by Captive Portal. This is useful for servers and devices without browsers.
- **Pass Listed Server Addresses**: Machines behind Captive Portal will be able to access these servers whether or not they have authenticated through Captive Portal. Typically these will be any DNS or DHCP servers that are separated from their clients by NG Firewall. If NG Firewall is handling DHCP or DNS, this is not necessary.

Figure 29: Passed Hosts Tab

Captive Page

This tab controls the type of Captive Portal page displayed to unauthenticated users. Please note that you can use HTML in the Captive Portal page fields, however invalid HTML will prevent the page from properly rendering.

- **Basic Message**: Select this option if users should see (or accept) a message before being allowed to the internet. It has several tunable properties such as **Page Title**, **Welcome Text**, **Message Text**, and **Lower Text**. If **Agree Checkbox** is enabled, users must check a checkbox (labeled with the **Agree Text**) before continuing.
- **Basic Login**: Select this option if users should see a page that requires them to login. Similar to **Basic Message**, it has several properties that can be configured. When the login/continue button on the page is clicked the user will be authenticated. You'll also need to set your authentication method on the **User Authentication** tab.
- **Custom**: Select this option if you would like to upload a custom Captive Portal page. This is for experienced web developers that are comfortable with HTML, Python and JavaScript - the Edge Threat Management Support department can not help with development or troubleshooting of custom Captive Portal pages. If **Custom** is selected it is advised to turn off automatic upgrades - newer versions of NG Firewall may be incompatible with any custom captive page.

> **Note:** When using 'Any OAuth provider' for User Authentication, you should select 'Basic Message'. All of the 'Page Configuration' options except for the agree checkbox and text will be used when generating the OAuth provider selection page.

Figure 30: Captive Page Tab

HTTPS/Root Certificate Detection

This feature checks if the root certificate is installed on the client machine. If the root certificate is not installed, you have the option to display a warning or block the connection. The Certificates used by HTTPS Inspector and other HTTPS connections to the unit including Captive Portal. This feature is highly recommended if you have HTTPS installed. The Certificates must have all the names and IP address used on the NG Firewall.

- **Disable Certificate Detection**: No checking for the root certificate.
- **Check Certificate. Show warning when not detected**: Checks the root certificate. If not found, displays a warning with instructions to install the certificate.

- **Require Certificate. Prohibit login when not detected**: Checks the root certificate. If the root certificate is not found, the connection is blocked and the client is given instructions to install the certificate.

The **Preview Captive Portal Page** button can be used to view what the configured captive page looks like. This button only works when Captive Portal in on.

Session Redirect

- **Block instead of capture and redirect unauthenticated HTTPS connections**: The browser redirecting from a HTTPS URL to the captive page will show certificate error as the captive page is not the page requested. To avoid this error message, block the traffic and show nothing instead of showing the captive login page.
- **Use hostname instead of IP address for the capture page redirect**: Create the browser redirect using the hostname instead of the IP address of the server.

> ⚠ **Warning:** If enabled, the admin must ensure that the hostname properly resolves to the internal IP of NG Firewall on all internal networks. If internal hosts use NG Firewall for DNS, this is automatic. If using another internal DNS server it is the administrator's responsibility to configure DNS to properly resolve to the correct internal IP on all internal networks. If this is not configured properly Captive Portal will not function properly as clients will not be able to reach the captive portal page. Host will **NOT** be able to reach the captive portal page if the hostname resolves to the external IP of NG Firewall.

This option is useful for organizations that have valid certificates on the NG Firewall server and wish to avoid the cert warning on the capture page.

> 📝 **Note:** This has nothing to do with the first warning caused by serving/spoofing the 301 redirect from a internet site to the capture page.

- **Always use HTTPS for the capture page redirect**: Always redirect to the HTTPS version of the login page when using Captive Portal.
- **Redirect URL**: Users will be rerouted to this site after successful authentication. If **Redirect URL** is blank they will be sent to the original destination. Make sure to enter a complete url (e.g. http://edge.arista.com) or this setting will not properly operate.

Custom Pages

You can create a custom.html file and place it along with any supporting image files, etc. into a zip file, and then upload the file via the administrative interface. This allows you to customize the look and layout of the page while leveraging the existing code and application settings. To use this model, you need to be familiar with HTML and forms.

Please note that customized Captive Portal pages are "use at your own risk". The Edge Threat Management Support department cannot assist you in creating, updating, or troubleshooting custom captive pages.

User Authentication

This section controls how users will be authenticated if the **Basic Login** page is used.

- **None**: is used in the case where no login is required.
- **Local Directory**: Use the NG Firewall's built-in Local Directory (**Config > Local Directory**) to authenticate users.
- **RADIUS**: Use an external RADIUS server to authenticate users. *This option requires Directory Connector to be installed and enabled and configured.*
- **Active Directory**: can be used if user should be authenticated against an Active Directory server. *This option requires Directory Connector to be installed and enabled and configured.*

- **Any Directory Connector**: can be used to allow users to authenticate against any of the configured and enabled Directory Connector methods. *This option requires Directory Connector to be installed and enabled and configured.*
- **Google Account**: can be used to allow users to authenticate via OAuth using a Google account.
- **Facebook Account**: can be used to allow users to authenticate via OAuth using a Facebook account.
- **Microsoft Account**: can be used to allow users to authenticate via OAuth using a Microsoft account.
- **Any OAuth Provider**: can be used to allow users to select and authenticate using any of the supported OAuth providers. When this option is selected, unauthenticated users will first encounter the OAuth selection page where they will click the icon or link corresponding to the provider account they wish to use.

The **Session Settings** section controls the timeout and concurrent login settings for Captive Portal.

- **Idle Timeout**: This option controls the amount of time before a host is automatically logged out if no traffic is seen. While a machine may be idle, it is still active on the network level. In this case **Idle** means no new TCP or UDP connections are seen by the Captive Portal.

  > ⚠️ **Important:** It is recommended to leave this at zero (not enabled).

- **Timeout**: This option controls the amount of time before a computer will be automatically logged out. After this the user must log in again through Captive Portal. Timeouts greater than 1440 minutes (1 day) is **not recommended**. The authenticated table is store in memory and will be flushed on reboot/upgrade. Additionally, the logout time should also be shorter than your DHCP lease time to assure IPs don't change before the Captive Portal timeout.
- **Allow Concurrent Logins**: This option controls if multiple machines can use the same login credentials simultaneously. If enabled, two or more users can login with the same username/ password at the same time.
- **Allow Cookie-based authentication**: When enabled, a cookie is added to the users browser and used to authenticate the user in future sessions. Cookies must be allowed by the browser and not cleared when closing the browser or by other security programs. When the Cookie timeout is reached the user is forced to re-authenticate (regardless of activity). The default is 24 hours.
- **Track logins using MAC address**: When enabled, Captive Portal will use the MAC address instead of IP address to identify the client machine. If the MAC address for a given IP address is not known it will revert to using an IP address. This option is useful on smaller flat networks where NG Firewall is on the same network segment as all the hosts, and you have a very long timeout period such that a client's IP address might change.
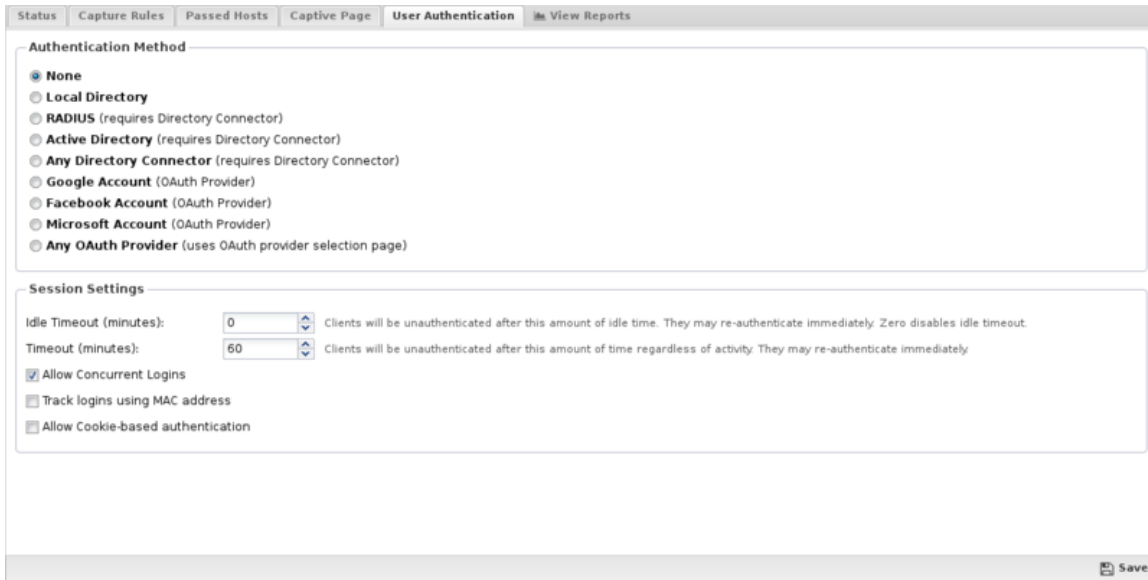
Figure 31: User Authentication Tab

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Captive Portal.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Captive Portal Summary	A summary of Captive Portal actions.
Activity Summary	A summary of Captive Portal activity.
Top Active Users	The top active users that logged in to Captive Portal.
Top Blocked Clients	The top clients that were blocked by Captive Portal because they were not logged in.
All Session Events	All sessions processed by Captive Portal.
Passed Session Events	Sessions matching passed hosts.
Captured Session Events	Sessions matching capture rules.
All User Events	All user sessions processed by Captive Portal.
Login Success User Events	Successful logins to Captive Portal.
Login Failure User Events	Failed logins to Captive Portal.
Session Timeout User Events	Sessions that reached the session timeout.
Idle Timeout User Events	Sessions that reached the idle timeout.
User Logout User Events	All user logout events.
Admin Logout User Events	Sessions logged off by the admin.

The tables queried to render these reports:

- Database Schema

Related Topics

Directory Connector

Captive Portal related topics

6.1.1   Captive Portal Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Captive Portal.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Captive Portal Summary	A summary of Captive Portal actions.
Activity Summary	A summary of Captive Portal activity.
Top Active Users	The top active users that logged in to Captive Portal.
Top Blocked Clients	The top clients that were blocked by Captive Portal because they were not logged in.
All Session Events	All sessions processed by Captive Portal.
Passed Session Events	Sessions matching passed hosts.
Captured Session Events	Sessions matching capture rules.
All User Events	All user sessions processed by Captive Portal.
Login Success User Events	Successful logins to Captive Portal.
Login Failure User Events	Failed logins to Captive Portal.
Session Timeout User Events	Sessions that reached the session timeout.
Idle Timeout User Events	Sessions that reached the idle timeout.
User Logout User Events	All user logout events.
Admin Logout User Events	Sessions logged off by the admin.

The tables queried to render these reports:

- Sessions
- Session Minutes

Related Topics

Report Viewer

Reports

6.2   IPsec VPN

The **IPsec VPN** service provides secure Internet Protocol (IP) communications by authenticating and encrypting each IP packet of a communication session.

Settings

This section reviews the different settings and configuration options available for IPsec VPN.

Status

The Status tab shows the status of the different components of the IPsec application.

- **Enabled IPsec Tunnels**

    This section shows a list of all IPsec tunnels that have been created and enabled. For tunnels that are active, the status will display the connection details reported by the IPsec subsystem. For inactive tunnels, the configuration information will be displayed.

- **Active VPN Sessions**

    This section shows a list of all active L2TP and Xauth connections. In addition to the connection details, there is a Disconnect column that can be used to forcefully disconnect an active session. Please note that there is no confirmation when you click the Disconnect icon. The corresponding session will be immediately terminated.

IPsec Options

- **Bypass all IPsec traffic**

    When this checkbox is enabled, traffic from IPsec tunnels will bypass all applications and services on the NG Firewall server. This was the only behavior available in previous versions of NG Firewall, so this option is enabled by default to maintain equivalent functionality on upgrade. If you disable this checkbox, traffic from IPsec tunnels can now be filtered through all active applications and services.

    Also please note that this only applies to plain IPsec tunnels. Traffic from L2TP and Xauth VPN clients will always pass through all active applications and services.

The IPsec Options tab contains common settings that apply to all IPsec traffic.

Unique ID's: Yes

The recommended default value is Yes unless you have a specific reason to change this setting.

Bypass all IPsec traffic: ☐

When enable, IPsec traffic will bypass processing by all Apps and Service Apps

💾 Save
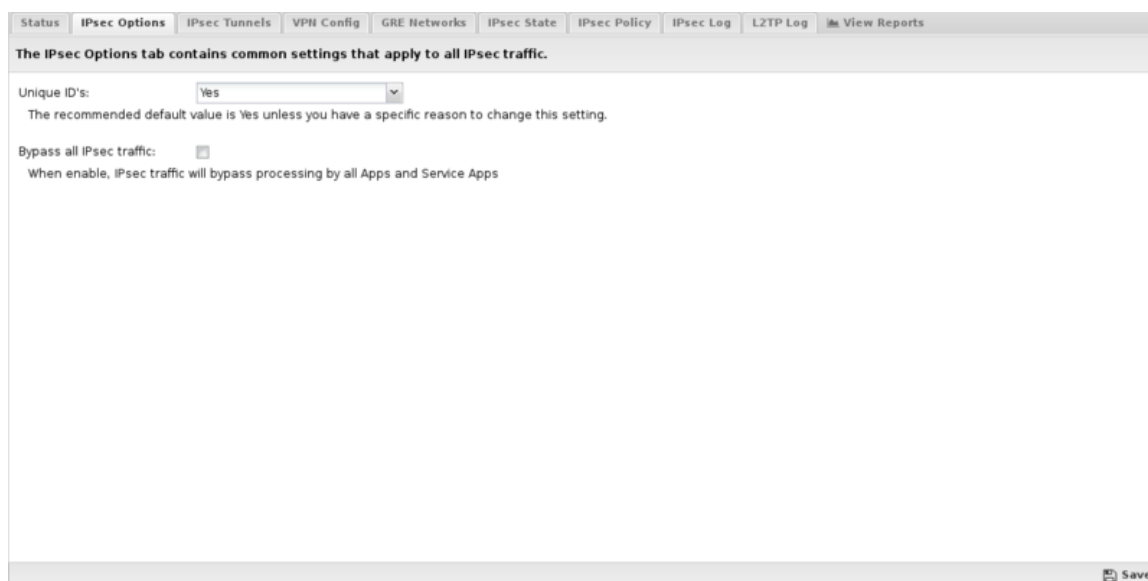
IPsec Tunnels

The IPsec Tunnels tab is where you create and manage the IPsec VPN configuration. The main tab display shows a summary of all IPsec tunnels that have been created.

- **Tunnel Editor**

  When you create a new tunnel, or edit and existing tunnel, the tunnel editor screen will appear with the following configurable settings:

Name	Description
Enabled	This checkbox allows you to set a tunnel to either enabled or disabled.
Description	This field should contain a short name or description.
Connection Type	This field allows you to set the connection type to any of the following: • Select Tunnel to specify a host-to-host, host-to-subnet, or subnet-to-subnet tunnel. This is by far the most common connection type. • Select Transport to specify a host-to-host transport mode tunnel. This connection type is much less common, and would generally only be used if you are attempting to establish an IPsec connection to another host which specifically requires this mode.
IKE Version	The IKE version to use, either version 1 or version 2. Both endpoints must use the same IKE version.

Name	Description
Connect Mode	This field controls how IPsec manages the corresponding tunnel when the IPsec process re-starts: • Select Always Connected to have the tunnel automatically loaded, routes inserted, and connection initiated. • Select On Demand to have the tunnel load in standby mode, waiting to respond to an incoming connection request.
Interface	This field allows you to select the network interface that should be associated with the IPsec tunnel on the NG Firewall server. For most situations choose **Active WAN** to bind to the active Internet interface. This allows the VPN tunnel to reconnect using a secondary WAN interface in the event of Internet failover. The **Active WAN** option is available when using the WAN Failover app. Alternatively, you can select a specific interface, or you can manually configure an IP address using the Custom option and manually inputting the IP address.
Any remote host	Enabling this option allows the VPN Server to accept tunnel connections from any IP Address. This option enables the remote side of the tunnel to connect from a dynamic IP address or via a secondary WAN link. This option switches the **Connect Mode** to **On Demand** and removes the Remote Host field as these options are not used when allowing connections from any remote host.
Remote Host	This field should contain the public IP address or DNS name of the host to which the IPsec VPN will be connected. ⚠️ **Warning:** Using host names with IPsec tunnels can often cause problems, especially if you have also enabled the L2TP/Xauth VPN server. We **strongly** recommend the use of IP addresses in the **Remote Host** field.
Local Identifier	This field is used to configure the local identifier used for authentication. When this field is blank the value in the *External IP* field will be used.
Remote Identifier	This field is used to configure the remote identifier used for authentication. When this field is blank, the value in the Remote Host field will be used. ❗ **Important:** If the remote host is located behind any kind of NAT device, you may need to use the value %any in this field for a connection to be successfully established.
Local Network	This field is used to configure the local network that will be reachable from hosts on the other side of the IPsec VPN.
Remote Network	This field is used to configure the remote network that will be reachable from hosts on the local side of the IPsec VPN.

Name	Description
Shared Secret	This field should contain the shared secret or PSK (pre-shared key) that is used to authenticate the connection, and must be the same on both sides of the tunnel for the connection to be successful. Because the PSK is actually used as the encryption key for the session, using long strings of a random nature will provide the highest level of security.
DPD Interval	The number of seconds between R_U_THERE messages. Enter *0* to disable this feature.
DPD Timeout	The number of seconds for a dead peer tunnel to be restarted.
Ping Address	The IP address of a host on the remote network to ping for verifying that the tunnel is connected and routing. Leave blank to disable.
Ping Interval	The time in minutes between ping attempts of the ping address. Leave as *0* to disable. Recommended value is *1* when using a Ping address.
Authentication and SA/Key Exchange	If you leave the Phase 1 and Phase 2 manual configuration checkboxes disabled, IPsec will attempt to automatically negotiate the encryption protocol with the remote peer when creating the tunnel. Given the number of different IPsec implementations and versions, as well as the overall complexity of the protocol, best results can often be achieved by enabling manual configuration of these two options, and selecting Encryption, Hash, DH Key Group, and Lifetime values that exactly match the settings configured on the peer device.

| Status | IPsec Options | IPsec Tunnels | VPN Config | GRE Networks | IPsec State | IPsec Policy | IPsec Log | L2TP Log | 📊 View Reports |

The IPsec Tunnels tab is used to create and manage IPsec tunnels

➕ Add   ⬇ Import   ⬆ Export

Enabled	External IP	Remote Host	Local Network	Remote Network	Description		Edit	Copy	Delete
☐	198.51.100.1	203.0.113.1	192.168.101.0/24	192.168.102.0/24	Example 1		✏	🗐	🗑
☐	198.51.100.1	203.0.113.1	10.10.0.0/16	10.20.0.0/16	Example 2		✏	🗐	🗑

💾 Save

VPN Config

The VPN Config tab manages the L2TP/Xauth/IKEv2 server configuration, enabling VPN client connections from remote desktops and mobile devices. IPsec is a preferred option because it uses native VPN software built into most systems, and therefore does not require installation of 3rd party software. When available, use IKEv2 type of VPN connections for optimal performance and compatibility.

- **Enable L2TP/Xauth/IKEv2 Server**

  Use this checkbox to enable or disable the L2TP/Xauth/IKEv2 server.

- **L2TP Address Pool**

  This field configures the pool of IP addresses that will be assigned to L2TP clients while they are connected to the server. The default *198.18.0.0/16* is a private network that is generally reserved for internal network testing. It was chosen as the default because it is used less frequently than other *RFC-1918* address blocks, and thus is less likely to conflict with existing address assignments on your network.

- **Xauth Address Pool**

  This field configures the pool of IP addresses that will be assigned to Xauth clients while they are connected to the server. The default *198.19.0.0/16* is a private network that is generally reserved for internal network testing. It was chosen as the default because it is used less frequently than other *RFC-1918* address blocks, and thus is less likely to conflict with existing address assignments on your network.

- **Custom DNS Servers**

  Leave both of these fields blank to have L2TP and Xauth clients use the NG Firewall server for all DNS resolution. Alternatively, if you have other DNS servers you want clients to use, you can enter IP addresses in these fields.

- **IPsec Secret**

  This is the shared secret that will be used between the client and server to establish the IPsec channel that will secure all L2TP and Xauth communications.

- **Allow Concurrent Logins**

  If enabled, the same credentials can be authenticated simultaneously from multiple devices.

- **User Authentication**

  In addition to the IPsec Secret configured above, VPN clients will also need to authenticate with a username and password. To use the Local Directory, select this option and click the **Configure Local Directory** button to manage use credentials. Alternatively, you can use an external RADIUS server for authentication by selecting the RADIUS option, and clicking the Configure RADIUS button to configure the RADIUS server options.

- **Server Listen Addresses**

  This list is used to configure one or more of your server IP addresses to listen for inbound VPN connection requests from remote clients. Clicking the add button will insert a new line allowing the entry of another server IP address.

The VPN Config tab contains settings used to configure the server to support IPsec L2TP, Xauth, and IKEv2 VPN client connections.

Enable L2TP/Xauth/IKEv2 Server: ☐

💾 Save

GRE Networks

The GRE Networks tab is where you create and manage connections to remote GRE servers. Generic Routing Encapsulation (GRE) is a tunneling protocol that can encapsulate a wide variety of network layer protocols inside virtual point-to-point links over an Internet Protocol network.

GRE Address Pool

This field configures the pool of IP addresses that will be assigned to interfaces created and associated with tunnels added on the GRE Networks tab. The default *198.51.100.0/24* is a private network that is generally reserved for internal network testing. It was chosen as the default because it is used less frequently than other *RFC-1918* address blocks, and thus is less likely to conflict with existing address assignments on your network. If you use GRE to connect multiple NG Firewall servers together, you need to configure a different, unused pool on each server.

The main tab display shows a summary of all GRE Networks that have been created.

Network Editor

When you create a new GRE Network, or edit and existing network, the network editor screen will appear with the following configurable settings:

Name	Description
Enable	This checkbox allows you to set a network to either enabled or disabled.
Description	This field should contain a short name or description.
Interface	This field allows you to select the network interface that should be associated with the GRE Network on the NG Firewall server. When you select a valid interface, the Local IP field (see below) will automatically be configured with the corresponding IP address. If for some reason you want to manually configure an IP address that is not currently active, you can set the Interface to Custom and manually input the IP address below.

Name	Description
External IP	Use this field to configure the IP address that is associated with the GRE Network on the NG Firewall server. Normally this field will be read-only and will automatically be populated based on the Interface selected above. If you select Custom as the interface, you can then manually enter the local IP address.
Remote Host	This field should contain the public IP address of the host to which the GRE tunnel will be connected.
Remote Networks	This field is used to configure the list of remote network traffic that should be routed across this GRE tunnel. Networks should be entered one per line in CIDR (***192.168.123.0/24***) format. **Note:** The subnets in Remote Networks must not conflict with your GRE Address Pool.

IPsec State

The IPsec State tab allows you to see the status of all established IPsec connections. There will typically be two entries per tunnel, one with details about the local side of the connection, and another with details about the remote side of the connection.
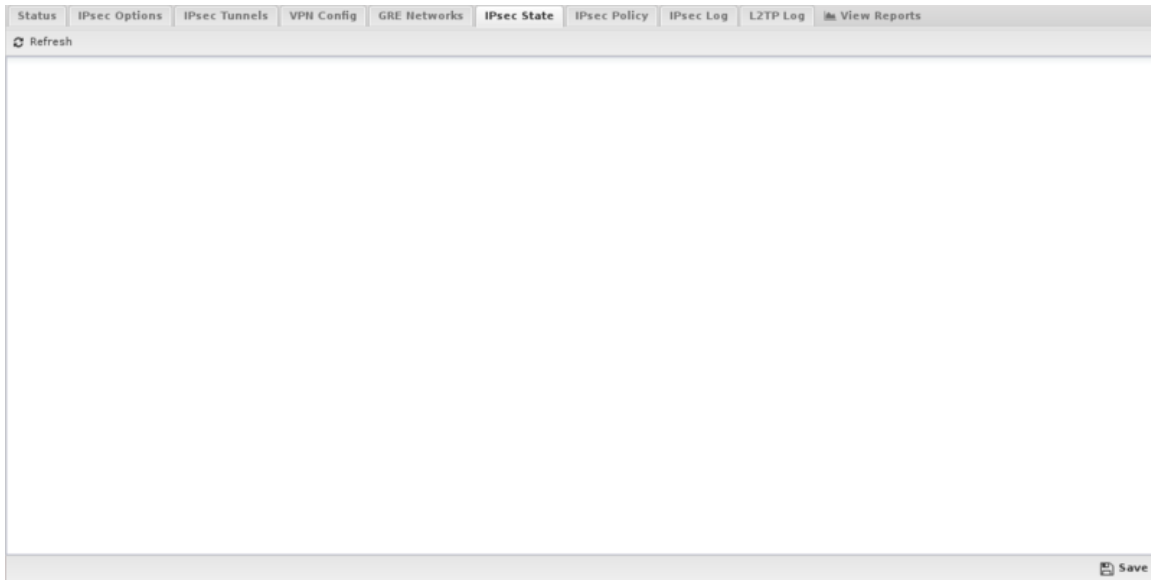
IPsec Policy

The IPsec Policy tab allows you to see the routing table rules associated with each IPsec VPN that is active.
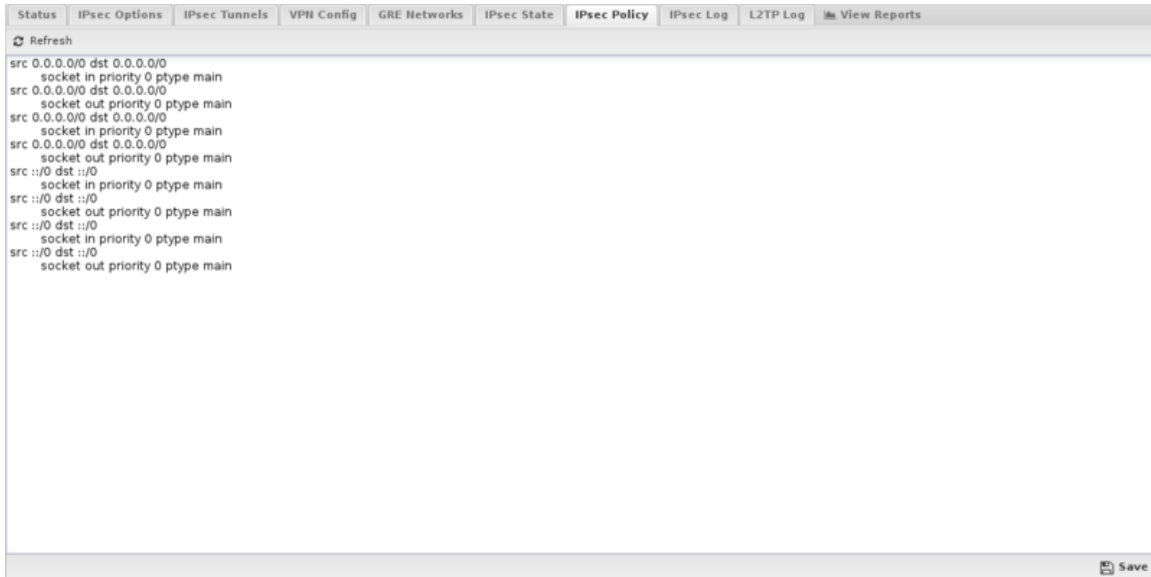
IPsec Log

The IPsec Log tab allows you to see the low level status messages that are generated by the underlying IPsec protocol components. This information can be very helpful when attempting to diagnose connection problems or other IPsec issues.
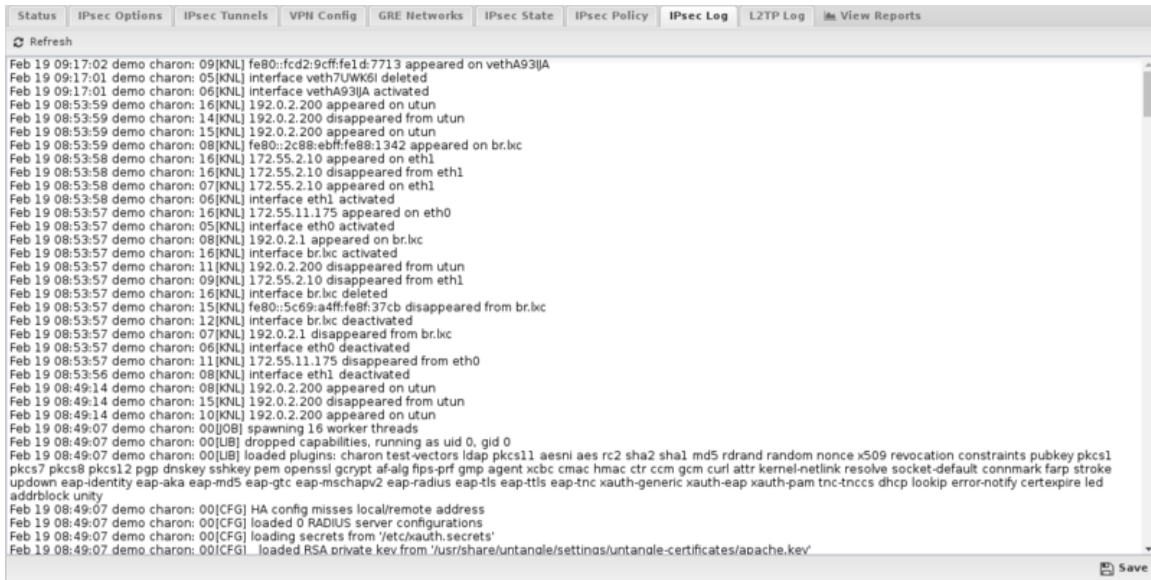
L2TP Log

The L2TP Log tab allows you to see the low level status messages that are generated by the underlying L2TP protocol daemon. This information can be very helpful when attempting to diagnose connection problems or other L2TP issues.

Reporting

The **Reports** tab provides a view of all reports and events for all connections handled by IPsec VPN.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
IPsec VPN Summary	A summary of IPsec VPN actions.
Hourly Tunnel Traffic	The amount of IPsec tunnel traffic over time.
Top Tunnel Traffic	The amount of traffic for each IPsec tunnel.
Top Active Users	The top IPsec VPN users by number of sessions.
Top Download Users	The top IPsec users grouped by amount of data downloaded.
Top Upload Users	The top IPsec users grouped by amount of data uploaded.
Top Protocols	The top IPsec VPN connections by protocol.
L2TP/Xauth Events	Shows all user L2TP/Xauth events.
Tunnel Connection Events	Shows all IPsec VPN tunnel connection events.
Tunnel Traffic Events	Shows all IPsec tunnel traffic statistics events.

The tables queried to render these reports:

- Database Schema
- Reporting

Related Topics

OpenVPN

6.2.1   IPsec VPN Reports

The **Reports** tab provides a view of all reports and events for all connections handled by IPsec VPN.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
IPsec VPN Summary	A summary of IPsec VPN actions.
Hourly Tunnel Traffic	The amount of IPsec tunnel traffic over time.
Top Tunnel Traffic	The amount of traffic for each IPsec tunnel.
Top Active Users	The top IPsec VPN users by number of sessions.
Top Download Users	The top IPsec users grouped by amount of data downloaded.
Top Upload Users	The top IPsec users grouped by amount of data uploaded.
Top Protocols	The top IPsec VPN connections by protocol.
L2TP/Xauth Events	Shows all user L2TP/Xauth events.
Tunnel Connection Events	Shows all IPsec VPN tunnel connection events.

Report Entry	Description
Tunnel Traffic Events	Shows all IPsec tunnel traffic statistics events.

The tables queried to render these reports:

- Database Schema

6.3   OpenVPN

OpenVPN enables you to create an SSL-based VPN (virtual private network) that supports both site-to-site and client-to-site tunnels. This allows your **road warrior** users to connect to local resources as if they were in the office, or connect the networks of several geographically distant offices together - all with the added security of encryption protecting your data. OpenVPN supports any operating system with an OpenVPN-compatible VPN client (which is almost every OS), even smartphones!The OpenVPN application can run as a server allowing for remote clients to connect to the NG Firewall server, and the OpenVPN application can connect to other remote NG Firewall servers as a client. The VPN Overview article provides some general guidance of which VPN technology may be the best fit for different scenarios.

Settings

This section reviews the different settings and configuration options available for OpenVPN.

Status

The **Status** tab shows you a list of open connections, the time the tunnels were created and transmit statistics.

Connected Remote Clients

This grid shows the currently connected remote clients that are connected to this OpenVPN (if server is enabled.)

Name	Description
Address	The IP of the remote client.
Client	The OpenVPN client name.
Start Time	The time that the client connected.
Rx Data	The amount of data received from this client in this session.
Tx Data	The amount of data sent to this client in this session.

Remote Server Status

This grid shows the remote servers this OpenVPN is connecting to as a client.

Name	Description
Name	The name of the remote server.
Connected	The current connection status
Rx Data	The amount of data received from this client in this session.
Tx Data	The amount of data sent to this client in this session.
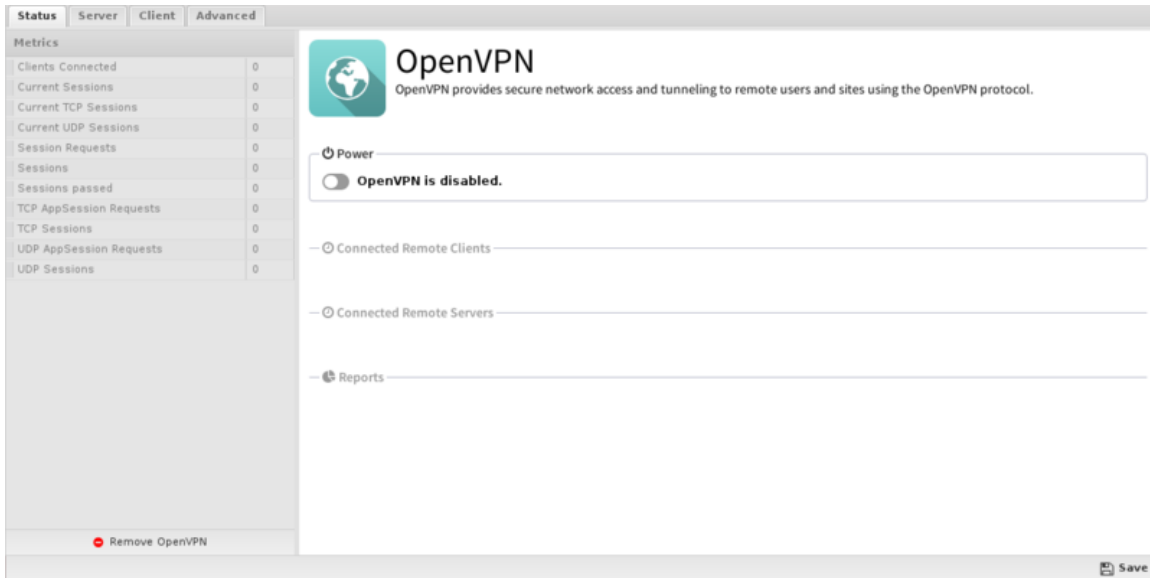
Server
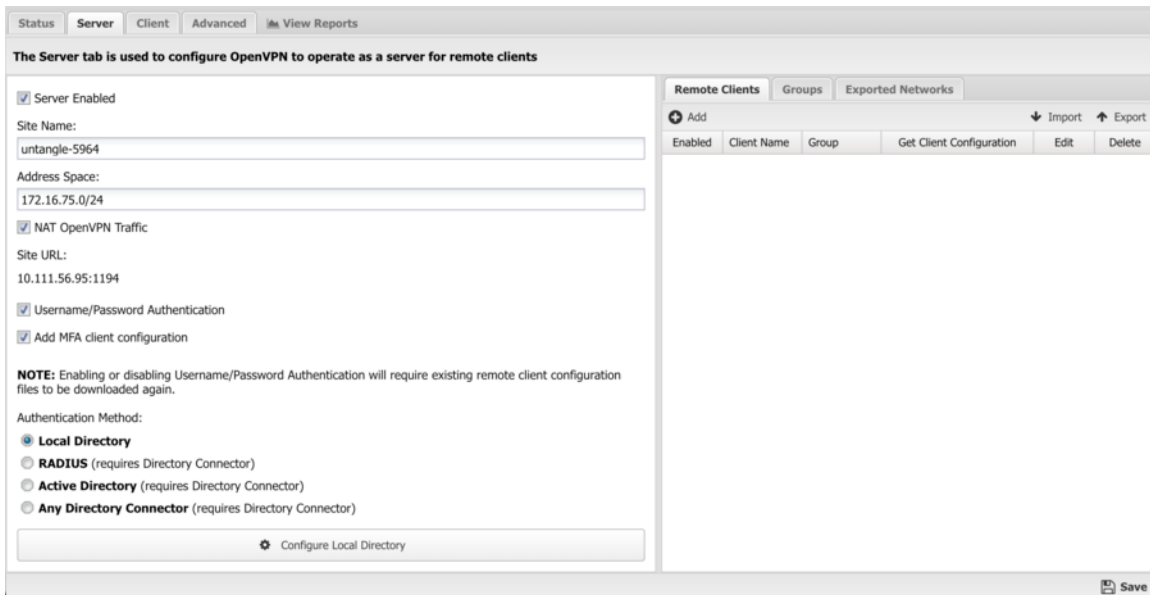
The Server tab includes all the configuration for OpenVPN's server functionality.

Site Name is the name of the this OpenVPN site. A random name is chosen so that it is unique. A new name can be given, but it should be unique across all NG Firewall sites in the organization. For example, if the company name is "MyCompany" then "mycompany" is a bad site name if you have multiple NG Firewalls deployed as it might be used elsewhere. The **Site Name must** be unique.

Site URL shows the URL that remote clients will use to connect to this server. This is just for reference. Verify that this address will resolve and be publicly reachable from remote networks. This URL can be configured in **Config** > **Network** > **Hostname**.

If **Server Enabled** is checked, the OpenVPN server will run and accept connections from configured **Remote Clients**. If unchecked the OpenVPN server will not run and not server services will be provided.

Address Space defines an IP network/space for the VPN to use internally. The **Address Space must** be unique and separate from all existing networks and other address spaces on other OpenVPNs. A default will be chosen that does not conflict with the existing configuration.

NAT OpenVPN Traffic will NAT all traffic from remote networks to local networks to a local address. This helps solve routing and host-based firewall issues. The default and recommended value is enabled.

Username/Password Authentication can be enabled to activate two factor authentication, requiring clients to also provide a username and password when connecting.

Add MFA client configuration can be enabled to activate multi-factor authentication using a TOTP app. This feature uses the Local Directory users and requires each user to be configured with multi-factor authentication and paired with a TOTP app.

Authentication Method is used to select the authentication method for clients when **Username/Password** authentication is enabled.

Remote Clients

The Remote Clients sub-tab configures all the **Remote Clients** that can connect to this OpenVPN server. A **Remote Client** is any entity that connects to this OpenVPN server as a client. This includes both remote desktops, laptops, devices, road warriors, etc. This also includes remote OpenVPNs and remote NG Firewall networks.

Initially there are no clients that are allowed to connect and an unique entry must be created for each remote client you wish to allow to connect to this server.

To add a new **Remote Client** click on the **Add** and provide the following information:

- **Enabled** - If checked, this client is enabled. If unchecked, this client is disabled and can not connect.
- **Client Name** - A unique name for the client. (alphanumerics only)
- **Group** - The **group** for this client. More information below.
- **Type** - The type of this client. **Individual Client** for a single host like a remote desktop or laptop. **Network** for an entire remote network that the server should also be able to reach.
- **Remote Networks** - The remote network in CIDR notation if this remote client is of type **Network**. For example: *192.168.1.0/24* means that the *192.168.1.** network lives behind the remote client and should be reachable from the server. If there are multiple networks reachable through this remote client, a comma separated list of CIDR networks can be used. These networks are automatically **exported** such that hosts on the main network and other remote clients can reach these networks.

After configuring this information save the new **Remote Client** by clicking **OK** then **Apply.** After saving settings, click on the **Download Client** button in the **Remote Clients** table on the row for the new client. This will provide links to download the configuration profile for the configured client.

- **Click here to download this client's configuration zip file for other OSs (Apple/Linux/etc).** provides a zip file with the OpenVPN client configuration files. This file can be used to configure various OpenVPN clients for various OSs, like linux, apple, and even some phones/tables/devices.
- **Click here to download this client's configuration file for remote OpenVPN clients.** provides a zip file with the OpenVPN client configuration for setting up a remote OpenVPN application to connect as a client to this server.
- **Click here to download this client's configuration onc file for Chromebook.** provides a onc file that can be used to configure your Chromebook as a client to connect to the NG Firewall OpenVPN server. On the target device, browse to chrome://net-internals and use Import ONC file.

On the client system, you must first install the OpenVPN client. You can download client from here: https://openvpn.net/download-open-vpn/. After installing the OpenVPN client on the remote client, you can import the OpenVPN profile into the client.

> **Note:** A client can only be connected once. If you install the same client on multiple remote devices they will kick each other off when a new one logs in. In most cases you need to setup a client for each remote device.

Groups

Groups are convenience feature to "group" clients together and apply some settings to that entire group. By default there will be a **Default Group.** Each group has the following settings:

- **Full Tunnel** - If checked, remote clients will send ALL traffic bound to the internet through the VPN. This allows for NG Firewall to filter ALL internet traffic for connected clients by "proxying" it through the VPN and then out through NG Firewall's internet connection. This will have no effect on remote OpenVPN clients. If unchecked, then only traffic destined to the local network are subject to filtering.
- **Push DNS** - If enable, OpenVPN will "push" some DNS configuration to the remote clients when they connect. This is useful if you wish for some local names and services to properly resolve via DNS that would not publicly resolve.
- **Push DNS Server** - If set to **OpenVPN Server** then the IP of the NG Firewall server itself will be pushed to the remote clients and all remote clients will use NG Firewall for all DNS lookups. If **Custom** is selected then one or two DNS entries can be specified that will be used for DNS resolution.
- **Push DNS Custom 1** - If **Push DNS Server** is set to custom, this IP will be pushed to remote clients to use for DNS resolution. It is important to export this address if that traffic should travel through the VPN tunnel. If this value is blank nothing will be pushed.
- **Push DNS Custom 2** - Just like **Push DNS Custom 1** except this sets the secondary DNS value. If blank, no secondary DNS will be pushed.
- **Push DNS Domain** - If set this domain will be pushed to remote clients to extend their domain search path during DNS resolution.

These settings will apply to all clients belonging to that group. Many sites will only have one group because all clients need the same settings. However, some clients have some **Full Tunnel** remote clients and some **Split Tunnel** remote clients. In this case, you need two groups where each client belongs to the appropriate group.

Exported Networks

Exported Networks is a list of networks that are reachable through the OpenVPN server for remote clients. **Exported Networks** are routes that are pushed to remote clients when they connect effectively telling remote clients to reach the specified network through the OpenVPN server.

For example, exporting *1.2.3.4/24* will result in all *1.2.3.* traffic going through the OpenVPN server.

The **Exported Networks** grid is pre-populated on installation with the IP/netmask of each static non-WAN interface.

- If **Enabled** is checked this network will be exported/pushed to connecting remote clients.
- **Export Name** is a name, purely used for documentation purposes.
- **Network** is the network, in CIDR notation.

Client

The **Client** tab is used to configure which remote servers this OpenVPN will connect to as a client.

Remote Servers

The **Remote Servers** grid lists the currently configured remote servers that OpenVPN is configured to connect.

To configure a new server to connect to, first login to the remote server and configure a new client as described above and click on the **Download Client**. After you have downloaded the distribution zip file return to this OpenVPN and click on the **Browse** button below the **Remote Servers** grid. Select the zip file downloaded from the OpenVPN server and then press **OK**. Next press the **Submit** button to upload the zip file to OpenVPN which will add a new entry into the **Remote Servers** grid based on the configuration in the submitted zip file.

If the remote server requires Username/Password authentication, you will have to edit the configuration, enabled the **Username/Password authentication** checkbox, and enter the username and password to be used when establishing the connection.

Once connected to a remote server, you will be able to reach their exported networks. They will also be able to reach the networks on this server specified as the **Remote Network** configuration.

> **Note:** Site to Site connections are not full-tunnel even if selected in the Group for the site to site. Internet traffic on the remote site will exit through its local gateway.
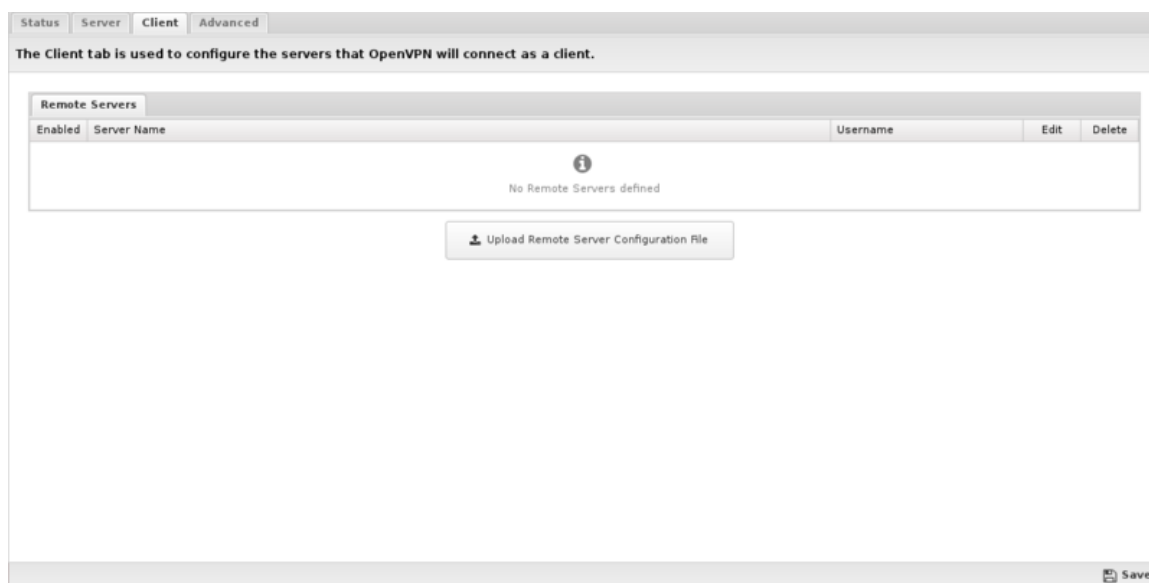
Advanced

The **Advanced** tab is provided for advanced users who have a detailed knowledge and understanding of OpenVPN, and need very specific configuration changes to address unique or unusual situations. It is entirely possible to completely break your OpenVPN configuration with a single wrong character, misplaced space, or by changing a configuration option that probably shouldn't be changed. Changes you make on this page can possibly compromise the security and proper operation of your sever, and are not officially supported.

Common Settings

At the top of the **Advanced** page are the Protocol, Port, and Cipher options. These must be the same on both the client and server for connections to work. Since they are the options most frequently modified, they can be easily configured here and will apply to both the client and server.

The **Client to Client Allowed** checkbox is used to enable or disable traffic passing between OpenVPN clients. When enable, all clients will have full network access to each other when connected. If disabled, traffic will not be allowed to flow between connected clients.

Server Configuration and Client Configuration

If you require changes to other low level parameters, the Server Configuration and Client Configuration grids allow you to effectively have total control of the OpenVPN configuration file that is generated. Both grids work the same way, with each configuration applied to the corresponding server or client openvpn.conf file respectively.

Both lists contain config items comprised of a Option Name and Option Value pair. By default, all items in both configuration grids are read only. The lists represent the default configuration settings used for the server and client configuration files. The default items cannot be modified or deleted, they can only be excluded. When you exclude an item, it is effectively removed from the resulting configuration file. To change one of the default items, simply add a new item with the same Option Name, and input the Option Value that you want to be used. This will effectively override the default. The same method can also be used to add configuration items that are not included in the default list.

Exclude Default Configuration Item

- This example shows how to disable the comp-lzo option in the server configuration file to turn off compression:

Modify Default Configuration Item

This example shows how to change the default keepalive setting in the server configuration file:

Add New Configuration Item

This example shows how to add a socks-proxy setting to the client configuration file:

Reports

The **Reports** tab provides a view of all reports and events for all connections handled by OpenVPN.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
OpenVPN Summary	A summary of OpenVPN actions.
OpenVPN Bandwidth Usage	The approximate amount of data transfered over openvpn connections.
OpenVPN Events	The amount of login and logout events over time.
OpenVPN Sessions	The amount of openvpn sessions over time.
Top Clients (by usage)	The number of bytes transferred grouped by remote client.
Connection Events	OpenVPN client connection events.
Statistic Events	Shows all OpenVPN connection traffic statistics events.

The tables queried to render these reports:

- Database Schema

Related Topics

OpenVPN

OpenVPN Reports

6.3.1   OpenVPN Reports

The **Reports** tab provides a view of all reports and events for all connections handled by OpenVPN.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
OpenVPN Summary	A summary of OpenVPN actions.
OpenVPN Bandwidth Usage	The approximate amount of data transfered over openvpn connections.
OpenVPN Events	The amount of login and logout events over time.
OpenVPN Sessions	The amount of openvpn sessions over time.
Top Clients (by usage)	The number of bytes transferred grouped by remote client.
Connection Events	OpenVPN client connection events.
Statistic Events	Shows all OpenVPN connection traffic statistics events.

The tables queried to render these reports:

- Database Schema
- OpenVPN

6.4   Tunnel VPN

The **Tunnel VPN** service app provides secure tunnels to remote servers and services and determines which traffic on the network goes through these tunnels.

The VPN Overview article provides some general guidance of which VPN technology may be the best fit for different scenarios.

Use Cases

Tunnel VPN is used in a wide variety of configurations. Some common scenarios are described below.

Branch Offices

Organizations with one or more small branch offices can use Tunnel VPN to send all internet-bound traffic at the remote small branch through the central site for security and filtering. This alleviates the need to actively manage the security and filtering configuration at the branch offices, and also allows for easier management at the central site as well as centralized monitoring and reporting.

Remote Security Services

There are many cloud-based security services or **Cloud Access Security Brokers** (CASB) that will enforce policy and security network traffic as it transits from the local infrastructure to the internet.

Tunnel VPN can be configured to send traffic, either in total or selectively, to the desired cloud services. For example, Tunnel VPN can send all port 25 (SMTP) through a specific tunnel to a cloud email archiving service. Alternatively you could send, DNS, web, or even all traffic through dedicated cloud services.

SD-WAN

SD-WAN (software-denfined networking) type deployments often have the need to maintain several tunnels to dedicated CASBs or internet "exit" points. Tunnel VPN allows you to maintain connections to several cloud exit points and prioritize the tunnels such that if one tunnel goes down the next available tunnel will be leveraged.

When combined with WAN Failover and WAN Balancer this provides an easy way to ensure the network is always online and the best possible tunnel is being used for connectivity, regardless of cloud services going up or down or individual ISPs or internet connections being up or down.

Privacy

Tunnel VPN can connect to other NG Firewall services or most Privacy VPN services (like NordVPN, ExpressVPN, Private Internet Access, and so on).

Many countries have imposed limits or monitoring on "forbidden" content. This can range from content expressing certain political views, information on historical events, region-locked content, unapproved types of entertainment, or copyrighted material. Also many locations do not have access to ISPs (or governments) that respect net-neutrality.

For these locations, Tunnel VPN can provide safe encrypted passage to a location that supports a freer internet and supports net neutrality. Rules can either statically determine what traffic goes through a tunnel (specific hosts or ports) or can dynamically shift which traffic uses the tunnel using tags. For example, a host can be switched to using a tunnel once Skype or Bittorrent usage is detected.

Settings

This section reviews the different settings and configuration options available for Tunnel VPN.

Status

The Status tab shows the on/off status of Tunnel VPN.

Tunnels

The Tunnels tab configures the encrypted tunnels to remote servers/services.

To add a new tunnel, simply click the **Add** button at the top.

- **Enabled** - If checked, this tunnel is enabled. If not enabled it will not connect and not be active.
- **Tunnel Name** - A unique name for the tunnel.
- **Provider** - this is the remote service/provider. Select the appropriate option for the remote service.
    - **Arista** - this is for connecting to a remote NG Firewall server.
    - **NordVPN** - this is for connecting to NordVPN at [nordvpn.com].
    - **ExpressVPN** - this is for connecting to ExpressVPN at [expressvpn.com].
    - **Custom zip file** - used for any remote service that supplies an zip file with an OpenVPN configuration inside.
    - **Custom zip file with username/password** - used for any remote service that supplies an zip file with an OpenVPN configuration inside and also requires a valid username and password.
    - **Custom ovpn file** - used for any remote service that supplies an ovpn file.

- **Custom ovpn file with username/password** - used for any remote service that supplies an ovpn file also requires a valid username and password.
  - **Custom conf file** - used for any remote service that supplies an openvpn conf file.
  - **Custom conf file with username/password** - used for any remote service that supplies an openvpn conf file also requires a valid username and password.
- **Select VPN Config File** button uploads the `zip/conf/ovpn` file.
- **Username** specifies the username (if required).
- **Password** specifies the password (if required).

First, provide a name and choose the remote provider type. After choosing the provider type the instructions will describe how to configure the rest of the fields.

On save, all enabled tunnels will attempt to connect to the remote services. The log can be viewed on the Log tab.

Rules

Rules control what traffic is routed through the tunnels. The Tunnel VPN rules are run before any WAN Balancer rules are evaluated and before the routing table is consulted. If a Tunnel VPN rule matches and the tunnel is active the traffic will exit through the tunnel regardless of the WAN Balancer or routing configuration. In other words Tunnel VPN takes precedence over any other routing configuration.

The Rules describes how rules work and how they are configured. As with all rules, rules are evaluated in order and the action is taken from the first matching rule.

Example: Static Rules

- If allow of the following conditions are met:
  - **Destination Port** is *25*
- Perform the following action(s):
  - **Destination Tunnel:** tunnel-1

This will route all **port 25** traffic through **tunnel-1**. If **tunnel-1** is offline, traffic will be routed normally.

Example: Preference Order

- **Rule 1**: Always (no conditons) perform the following action, Destination Tunnel: '**tunnel-1**'
- **Rule 2**: Always (no conditons) perform the following action, Destination Tunnel: '**tunnel-2**'
- **Rule 3**: Always (no conditons) perform the following action, Destination Tunnel: '**tunnel-3**'

Then traffic will always route to **tunnel-1**. If **tunnel-1** is not available it will route to **tunnel-2**. If **tunnel-2** is not available it will route to **tunnel-3**. If **tunnel-3** is not available it will route normally.

Example: Dynamic Rules

Unlike most solutions, NG Firewall also allows for automatic dynamic adjustment of what traffic goes through the tunnel by using tags. Hosts can be tagged manually by tagging the appropriate device or username associated with a host, or automatically using trigger rules Events.

For example if you'd like a host using bittorrent to automatically be routed through the tunnel. Add a trigger rule to tag hosts detected as using bittorrent (an example is there by default) and then add the following Tunnel VPN rule:

- If allow of the following conditions are met:
  - **Client Tagged** is *bittorrent-use*.
- Perform the following action(s):
  1. **Destination Tunnel: *tunnel-1***.

This will route any hosts tagged **bittorrent-use** through **tunnel-1**. The trigger rule will ensure that any host detected using Bittorrent will automatically be tagged so that each session after the detection will go through the tunnel.

Example: Multiple Triggers

If there are many scenarios in which a host should be routed through a tunnel you can configured multiple triggers. For example you can configure multiple trigger rules:

- If host is using Skype, tag host **tunnel** expires in **10** minutes.
- If host is accessing craigslist, tag host **tunnel** expires in **10** minutes.
- If host is accessing **Gaming** category website, tag host **tunnel** expires in **10** minutes.

Then add the following Tunnel VPN Rule:

- If allow of the following conditions are met:
    - **Client Tagged** is **tunnel**.
- Perform the following action(s):
    - **Destination Tunnel: tunnel-1**.

And if a host does any of those action, it will automatically be switched to the **tunnel** (until the tag expires which will be **10** minutes after the spefied activity stops).

Log

This shows the raw OpenVPN log file. Beware: there are often many errors logged by OpenVPN that are not an issue.

This is useful for debugging issues if the tunnels are not initializing correctly to the service providers.

Reports

There are currently no specific reports for Tunnel VPN. However all traffic is logged with the appropriate tunnel set as the destination interface.

All reports (Application Control, Web Filter etc.) can be viewed and filter per tunnel by adding a **Destination Interface** condition where the value equals the tunnel ID.

Related Topics

OpenVPN

OpenVPN Reports

6.4.1   Tunnel VPN Reports

There are currently no specific reports for Tunnel VPN. However all traffic is logged with the appropriate tunnel set as the destination interface.

All reports (Application Control, Web Filter, etc.) can be viewed and filter per tunnel by adding a *Destination Interface* condition where the value equals the tunnel ID.

6.5   WireGuard VPN

The **WireGuard VPN** service provides virtual private networking via Wireguard VPN, which is an open source lightweight VPN application and protocol designed to be fast, secure, and easy to configure.

Settings

This section reviews the different settings and configuration options available for WireGuard VPN.

Status

The Status tab shows the status of the WireGuard VPN service

• **Local Service Information**

> This section displays information about the local WireGuard service such as the public key, endpoint address and port, peer address, and the list of local networks.

• **Enabled Tunnels**

> This section shows a list of active WireGuard tunnels.

Settings

• **Listen port**

> Sets the port where the WireGuard server will listen for inbound tunnel connections from peers.

• **Keepalive interval**

> Sets the passive keepalive interval which ensures that sessions stay active and allows both peers to passively determine if a connection has failed or been disconnected.

• **MTU**

> Sets the MTU size for WireGuard tunnels.

Remote Client Configuration

These fields are used when generating the Remote Client configuration.

• **DNS Server**

IP Address of local DNS server that will be added to client configuration. It is initially populated using the first defined DHCP DNS Server Override address is used it found. If not, the IP address of your first non-WAN interface is used.

- **Networks**

    These are networks added to the client's allowed IP list. It is initially populated with all known local networks discovered from non-WAN interfaces (and their aliases) and static routes.

Peer IP Address Pool

- **Assignment**

    Used to select the method for address pool assignment. Can be set for **Automatic** to allow the system to automatically select an unused network space or **Self-assigned** to configure a user entered network space.

- **Network Space**

    Shows the automatically assigned networks space or allows editing the self-assigned network space.

- **New Network Space**

    Click when using Automatic assignment to select a new random network space.

Tunnels

The Tunnels tab is where you create and manage WireGuard VPN tunnels. Each tunnel in the table has options to view the client configuration or edit the tunnel.

For a step by step guide to setting up WireGuard VPN tunnels, see Setting up WireGuard VPN site-to-site connections in NG Firewall.

- **Remote Client**

Clicking this icon will display a window showing the recommended client configuration in both Quick Reference (QR) Code, which many WireGuard mobile apps can scan with the devices camera and import and text file suitable for copying and pasting into the remote client.

- **Tunnel Editor**

When you add a tunnel, or edit and existing tunnel, the tunnel editor screen will appear with the following configurable settings:

Note: You can copy the configuration from a remote NG Firewall peer and paste it into any of the configurable fields. The screen automatically populates all of the relevant fields from the remote side. This simplifies the configuration of tunnels and is recommended to avoid misconfiguration.

Name	Description
Enabled	This checkbox allows you to set a tunnel to either enabled or disabled.
Description	This field should contain a short name or description.
Remote Public Key	This field is for the public key of the tunnel peer.
Remote Endpoint Type	This field controls the endpoint type for the peer. • Select **Roaming** if the remote endpoint is a mobile device using the WireGuard app, or if the remote network is used for client access only and does not host any resources. • Select **Static** for a traditional site to site tunnel configuration where each network hosts resources that must be accessible over the virtual private network.
Remote Endpoint IP Address	Sets the IP address for a static endpoint.
Remote Endpoint Port	Sets the port for a static endpoint.
Remote Peer IP Address	This field sets the IP address that will be used by the remote peer.
Remote Networks	This field is used to configure the list of remote networks that should be routed across this WireGuard tunnel. Networks should be entered on per line in CIDR (*192.168.123.0/24*) format.
Monitor Ping IP Address	The IP address of a host on the remote network to ping for verifying that the tunnel is connected. Leave blank to disable.
Monitor Ping Interval	The time in seconds between attempts to ping the configured ping monitor address.
Monitor Alert on Tunnel Up/ Down	When enabled, CONNECT and DISCONNECT alerts will be generated when the configured ping monitor transitions from reachable to unreachable, and unreachable to reachable.
Monitor Alert on Ping Unreachable	When enabled, UNREACHABLE alerts will be generated for each monitor ping that fails when the target is unreachable.

Name	Description
Local Service Information	This section includes information from the Status tab that is useful when doing copy/paste configuration between to peers.
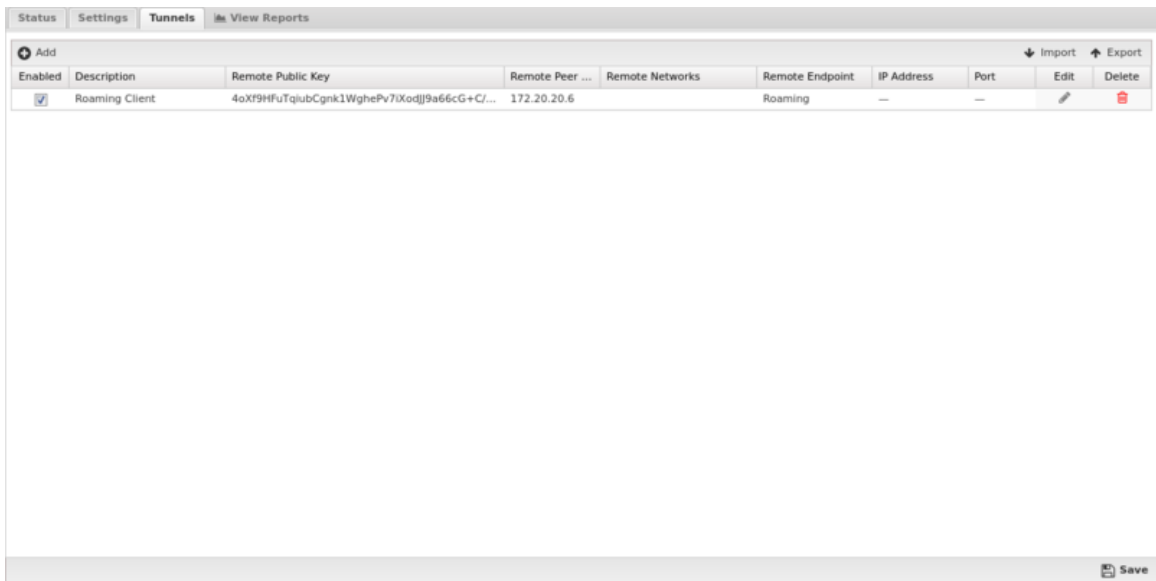
WireGuard VPN client

The WireGuard VPN client app is available for download on a variety of mobile device and desktop operating systems including iOS, macOS, Android, Windows, and Linux.The download links for each supported OS are available from the WireGuard Website.

For a step by step setup guide refer to the KB article Setting up WireGuard VPN on mobile devices and desktops.

Reporting

The Reports tab provides a view of all reports and events for all connections handled by WireGuard VPN.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
WireGuard VPN Summary	A summary of WireGuard VPN traffic.
WireGuard VPN Bandwidth Usage	The amount of traffic processed by the WireGuard service.
WireGuard VPN Events	Time chart of WireGuard VPN connection events.

Report Entry	Description
Top Remove Clients (by usage)	The top WireGuard VPN peers by traffic usage.
Connection Events	Shows all WireGuard VPN tunnel monitoring events.
Tunnel Traffic Events	Shows all WireGuard tunnel traffic statistics events.

The tables queried to render these reports:

- WireGuard VPN Reports

Related Topics

IPsec VPN

OpenVPN

6.5.1   WireGuard VPN Reports

The **Reports** tab provides a view of all reports and events for all connections handled by WireGuard VPN.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
WireGuard VPN Summary	A summary of WireGuard VPN traffic.
WireGuard VPN Bandwidth Usage	The amount of traffic processed by the WireGuard service.
WireGuard VPN Events	Time chart of WireGuard VPN connection events.
Top Remove Clients (by usage)	The top WireGuard VPN peers by traffic usage.
Connection Events	Shows all WireGuard VPN tunnel monitoring events.
Tunnel Traffic Events	Shows all WireGuard tunnel traffic statistics events.

NG Firewall Manage Apps

Contents

- Directory Connector
- Reports
- Policy Manager

7.1   Directory Connector

Directory Connector provides functionality to integrate with Microsoft's Active Directory or servers that support RADIUS Server, as well as some tools for managing the Host Viewer username mapping for the hosts on the network.

Figure 32: Directory Connector

Contents

- About Directory Connector
- Settings
  - Status
  - User Notification API
  - Active Directory
  - Azure Active Directory
  - RADIUS
  - Facebook
- Reports
- Related Topics

About Directory Connector

Directory Connector provides many tools to assist with Users.

Settings

This section reviews the different settings and configuration options available for Directory Connector.

Status

This displays the current status and some statistics.
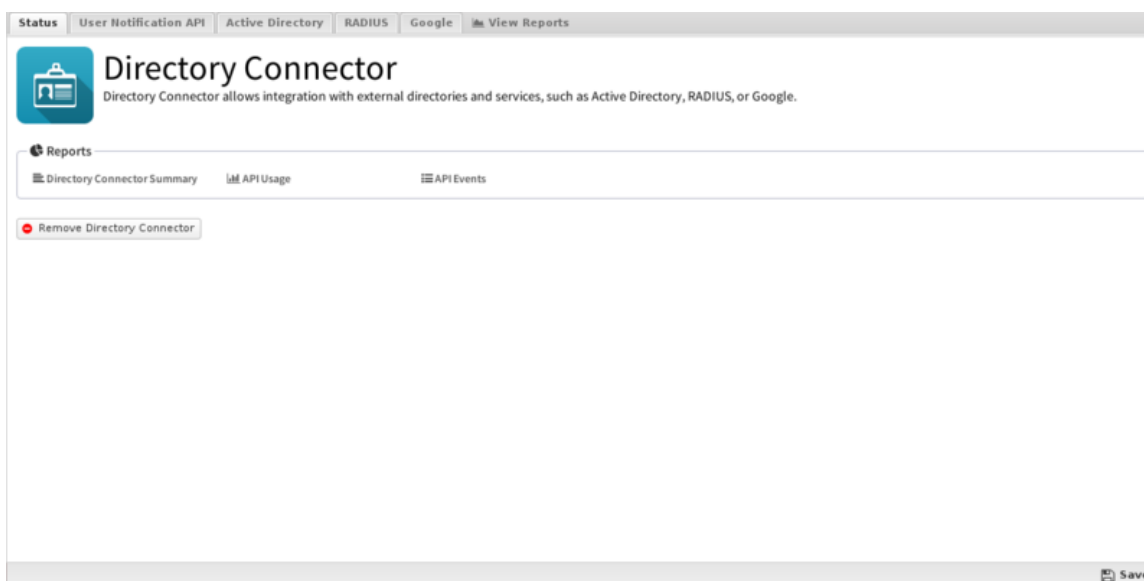
Figure 33: Directory Connector Status

User Notification API

The "User Notification API" is a webapp running on the NGFW that various external scripts can call to notify NG Firewall that a specific user is logged into a specific IP. The userapi webapp is used to update and maintain the associated usernames in the Hosts so that User Matcher in Rules match correctly. When a username is associated with the *Username* in Rules matches as expected.

This API can be called:

1. manually
2. via the *User Notification Login Script*
3. via the *Active Directory Server Login Monitor Agent*
4. via any custom script or external program

- **Enable/Disable** If enabled the User Notification API is enabled. If disabled, the User Notification is completely disabled.
- **Secret Key**: If specified, only API calls specifying the correct secret key will be allowed. All other requests are ignored. If not specified, it is not required to use the API however the clientIP argument is ignored to avoid API abuse.

The webapp lives at http://SERVERIP/userapi/registration on the server and can be called with the following arguments:

Argument	Example	Description
clientIp	*192.168.1.100*	The client IP address of the host in question
username	foobar	The username to associate with the client IP.
hostname	machinename	The hostname to associate with the client IP. This argument is optional.
action	*login* or *logout*	The action, *login* is assumed if no action is specified. *login* with associate the username and hostname of the specified client IP. *logout* will unset the client IP's associated username.
secretKey	foobarsecret	If this argument does not match the specified secretKey the call will be ignored.

For example, If the NGFW internal IP is **192.168.1.1** without a secretKey,
to associate user **foobar** on machine **foobarpc** to **192.168.1.100**
you would call visit this URL: http://192.168.1.1/userapi/registration?
action=login&clientIP=192.168.1.100&username=foobar&hostname=foobarpc&secretKey=foobarsecret.

To unset that username mapping when the client logs out simply visit this URL: http://192.168.1.1/
userapi/registration?action=logout&clientIP=192.168.1.100&secretKey=foobarsecret.

Obviously visiting these URLs manually each time a user logs in or out of a machine is not realistic.
Typically this process is automated in one of two ways described below or using a custom script.
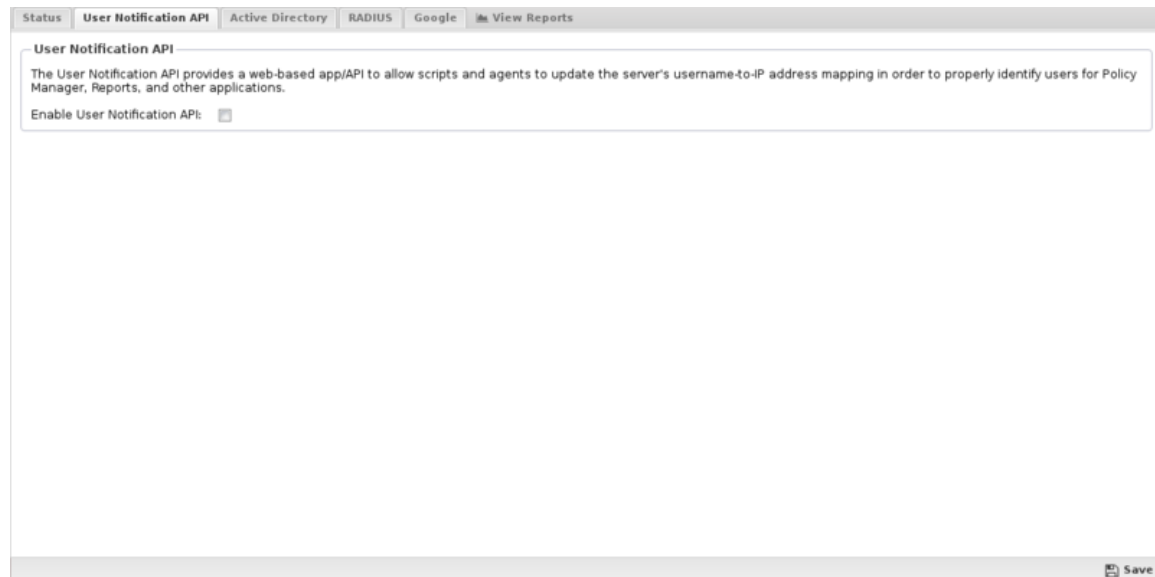
Figure 34: User Notification API tab

User Notification Login Script

The *User Notification Login Script* or *UNLS* which is a small script that runs at login on each machine
to notify the NGFW when a user logs in. This script can be pushed out to all the machines in a domain
via a group policy object. This is useful in cases where you want to set the username in the Hosts
without having users manually log into the Captive Portal.

Once installed, the script starts each time a user logs on to the network and immediately notifies
NG Firewall of the username and IP address. Once this process is finished, any activity for that IP
address will be automatically mapped to the username. This scripts runs on login and periodically in
the background to keep the Directory Connector Username Map updated with any current information
on your network users.

To download the User Notification Login Script, click on the **Download User Notification Login
Script** button and download the script. The script will be configured for your environment but may
require further customization. Review the script and make changes as needed.

Now that you have the UNLS on your Domain Controller, you need to decide if you want it run for
UNLS for the Entire Domain or UNLS for Specific Users.

UNLS for the Entire Domain

To apply UNLS to the your entire domain you'll need to set up a new Group Policy Object - please
follow the instructions below.

1. Click on the **Download User Notification Login Script** and save the `user_notification.vbs`
   file to `\\localhost\\NETLOGON`.

2. Log on to the Domain Controller, then launch the Group Policy Management Console (**Start > Run**: gpmc.msc).
3. From the Group Policy Management Console, right-click on the domain and select **Create and Link a GPO here**.
4. Specify a name for the Group Policy.
5. Right-click on the group policy that you just created and click **Edit**.
6. Go to **User Configuration > Windows Settings > Scripts (Logon/Logoff)**.
7. Click on the **Logon** icon, then **Show Files**. Windows Explorer will launch into the correct directory.
8. Copy the user_notification.vbs file that you downloaded to this location.
9. Click the **Add** button, browse for the script, then click **OK**.
10. In the Logon Properties window, click **Add** , type a descriptive script name, then click **OK**.
11. In the **Select User, Computer or Group** window, select the OU or Group to which you want to apply this GPO.
12. From a command prompt, activate the group policy that you just created: ***gpupdate /force***.

You can verify it is working by looking in the Event Log for login/logout events.

UNLS for Specific Users

If you only want to use the UNLS for a few users, you can use these instructions:

1. Click on the *Download User Notification Login Script* and save the `user_notification.vbs` file to `\\localhost\\NETLOGON`.
2. Using a text editor, create a `local.bat` file that has the following lines:

```
@ echo off
\\ADServerIPAddress\netlogon\user_notification.vbs
```

1. Save the `local.bat` file to `\\localhost\\NETLOGON`.
2. From the domain, go to the **Users** folder, right-click the user and go to **Properties**.
3. On the Profile tab, type the filename of the UNLS (probably `user_notification.vbs`) in the Logon script field.
4. Launch the Group Policy Management Console, then launch the Group Policy Object Editor (**Start > Run**: gpedit.msc).
5. Copy the `user_notification.vbs` file that you downloaded in the first step to this location.

Active Directory Server Login Monitor Agent

The other way to call the User Notification API is by running an agent/monitor on the Active Directory Server. The agent monitors the server's login events and updates the NG Firewall when a user logs to a computer. This has several advantages over the UNLS.

> **Note:** To use the Active Directory Login Monitor a Secret Key must be specified.

1. It allows you to set a secretKey that only the agent knows, so only the AD server itself can update the username mapping. (users have no way of overriding changing the information)
2. It is not necessary to run a login/logout script on all machines. No GPO is necessary.

First download and install the agent on the Active Directory server. and configure it so that it updates the NG Firewall server when it sees user login events.

Installation Guide

Download

Configure the *NGFW Settings* in Login Monitor so it updates your NG Firewall event when login events occur.
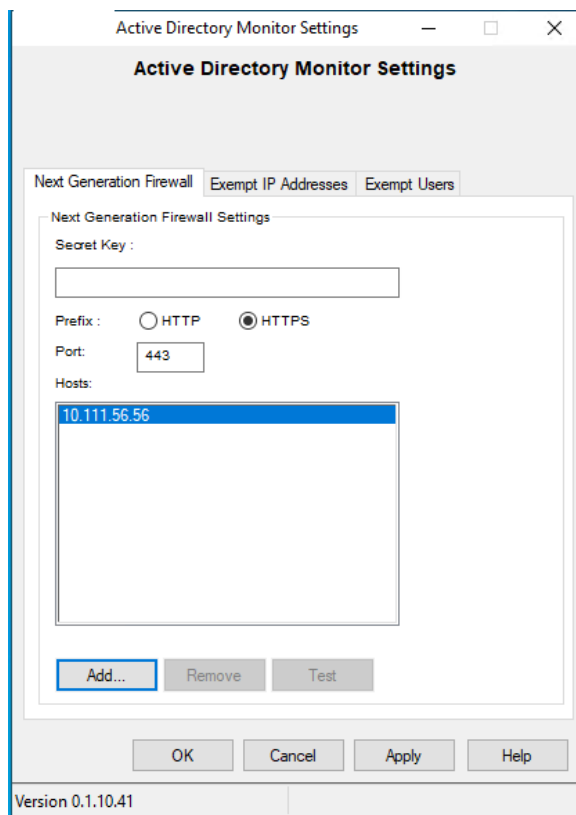
Figure 35: AD Server Login Monitor Settings

- **Secret Key**: The Secret Key if there is a *Secret Key* configured on the NGFW User Notification API. User Notification must be enabled on the NGFW. If no *Secret Key* is configured leave it blank.
- **Prefix**: The protocol to use to communicate with the NG Firewall.
- **Port**: The port to use to communicate with the NG Firewall. The default is port *80* for HTTP and *443* for HTTPS.
- **IP Addresses**: The IP addresses to reach your NG Firewall. Generally this should be the LAN addresses of your NG Firewall. By default HTTP and HTTPS is closed on the WAN side of NG Firewall. If the Login Monitor Agent cannot reach the NG Firewall, an error icon is shown next to the NG Firewall IP address entry.
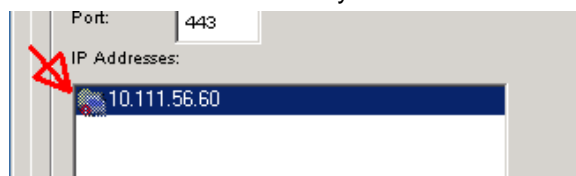
Figure 36: Error Reaching NG Firewall

The Exempt IP Addresses tab is a list of IP addresses which Login Monitor should ignore for login events. IP addresses are accepted in the following format:

- Single IP address (*192.168.2.2*)
- Wildcard IP address (*192.168.3.**)
- CIDR (*192.168.4.0/24*)
- Range (*192.168.5.5*-*192.168.5.102*)
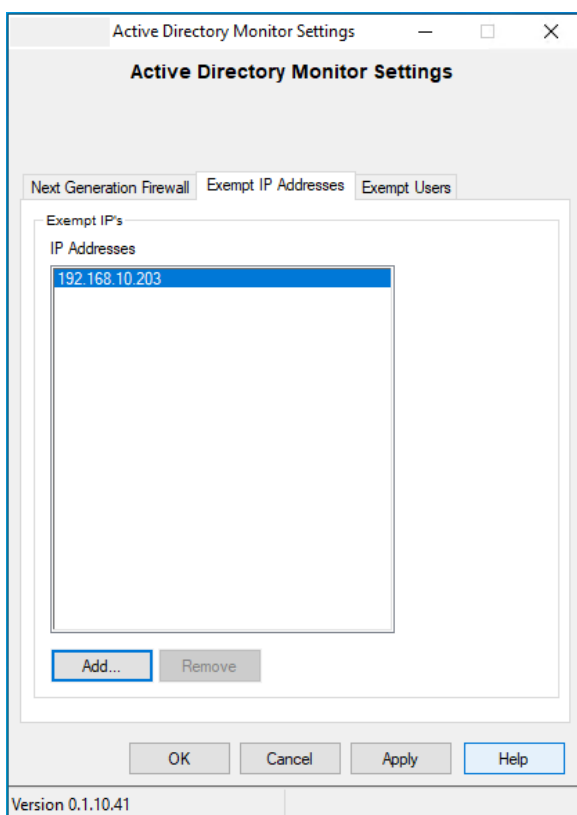
Figure 37: Exempt IP Address Tab

The Exempt Users tab is a list of AD users which Login show ignore for login events.

Active Directory

The Active Directory Connector allows NG Firewall to communicate with the Active Directory server. This is useful for two things:

1. Allowing users to login to Captive Portal using their AD login/password. The Captive Portal will verify the authentication information directly with the AD server.
2. Allow NG Firewall to query the groups so the it knows which groups a user belongs to. If this is configured the rules matcher in Rules will correctly match.

Before configuring the *Active Directory Connector* here are a few important steps:

1. Ensure that your Active Directory users are in one domain. Users can be in multiple Active Directory Organizational Units (OUs), but must be under one domain - multiple domains are not supported at this time.
2. Check to see if you have the Group Policy Management Console installed; if not, install it.

The Active Directory Connector tab contains settings for connecting and communicating with a Domain Controller. Other applications such as Captive Portal can use Directory Connector to authenticate and identify users against an existing Domain Controller.

- **AD Server IP or Hostname**: The IP or hostname of the AD server - we recommend using the IP to prevent DNS issues.
- **Secure**: Enable SSL for the connection to the AD server.
- **Port**: The port to use when connecting to the AD server. The default is 389.
- **Authentication Login**: Enter an Active Directory Administrator login.
- **Authentication Password**: Enter an Active Directory Administrator password.
- **Active Directory Domain**: Your domain, (e.g. mycompany.local).

- **Active Directory Organization**: The Active Directory organization unit (OU) that contains the users. If you want the NG Firewall server to find all users, leave this blank.

  If for some reason you want to limit the users to a specific part of the domain tree, specify the OU path in the format of OU=ouName. Only one OU can be entered.

You can use the test tools to verify your settings and view an *incomplete* user list. After Active Directory is configured, you can configure Captive Portal to use it for authenticating users if you wish.
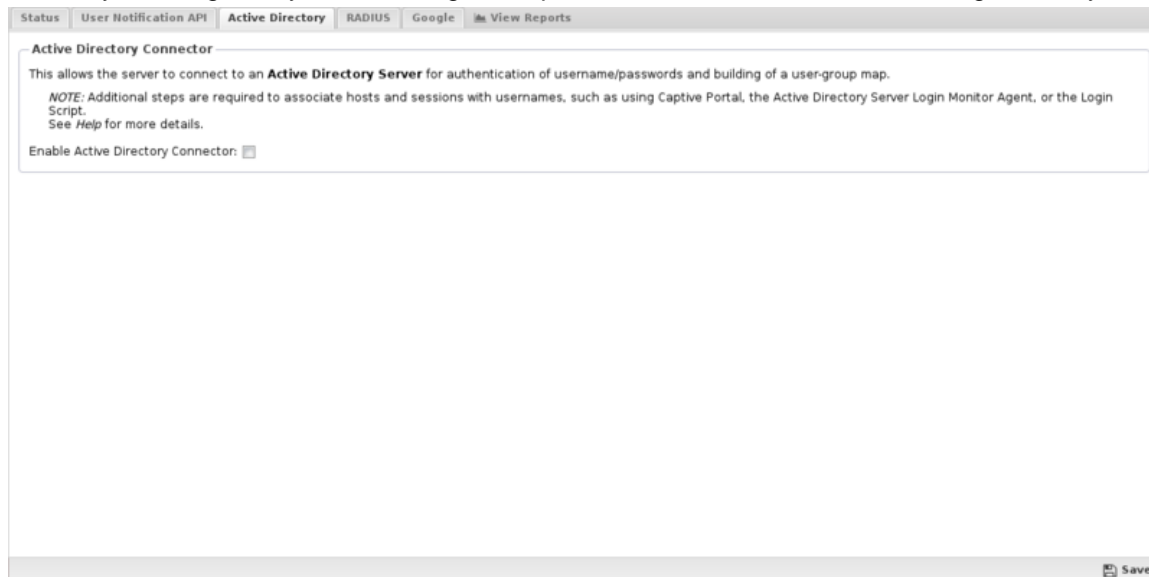
Figure 38: Active Directory tab

Azure Active Directory

You can use the Active Directory Connector to authenticate users against Azure Active Directory Domain Services. This type of connection requires that you have a Microsoft Azure account using Azure AD Domain Services. Before configuring NG Firewall to authenticate to your instance of Azure Active Directory, follow these steps:

1. Enable Azure Active Directory Domain Services
2. Generate an SSL certificate
3. Enable secure LDAP
4. Permit access to secure LDAP
5. Configure DNS for your AD domain

After you complete setting up Azure Active Directory, you can configure NG Firewall to authenticate via secure LDAP. The configuration of the connection is similar to a local Active Directory Domain except that you must enable the **Azure** checkbox. Confirm that the port is 636 and that **Secure** is enabled as Azure Active Directory requires secure LDAP.

RADIUS

The RADIUS Connector allows NG Firewall to communicate with a RADIUS server. This is useful for:

1. Allowing users to login to Captive Portal using their RADIUS login/password. The Captive Portal will verify the authentication information directly with the AD server.

The RADIUS tab contains settings to configure communication with the RADIUS server.

- **RADIUS Server IP or Hostname**: The IP or hostname of the RADIUS server - we recommend using the IP to prevent DNS issues.
- **Port**: The port to use when connecting to the RADIUS server. The default is *1812*.

- **Shared Secret**: This must match the shared secret set on the RADIUS server.
- **Authentication Method**: This must match the authentication method used by the RADIUS server.

You can use the test tool to verify your settings. After RADIUS is configured, you can configure Captive Portal to use it for authenticating users if you wish.
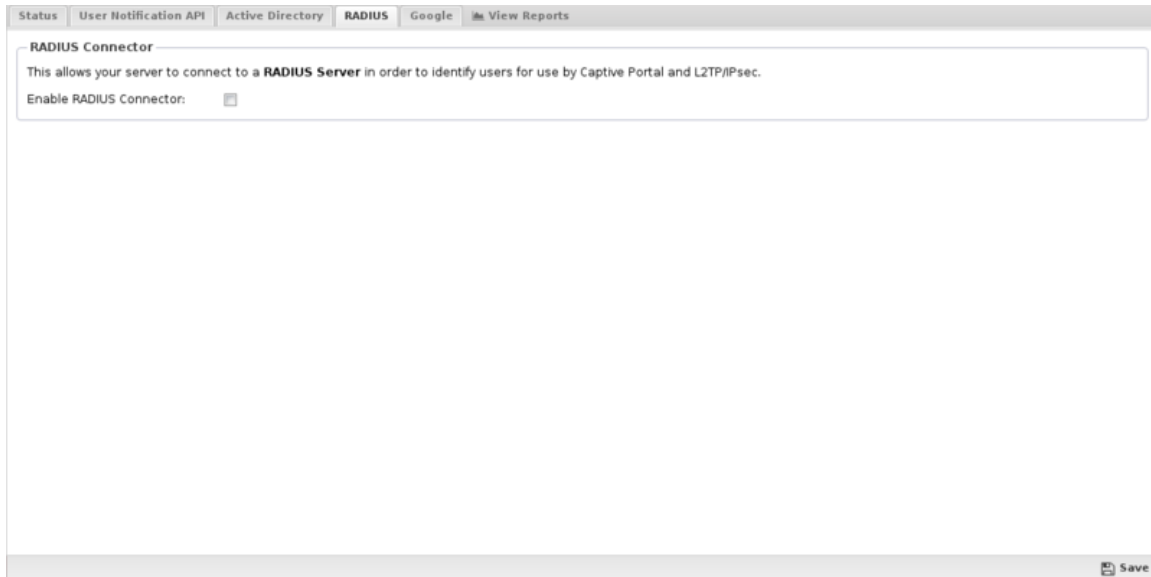
Figure 39: RADIUS Connector

Facebook

The Facebook Connector allows NG Firewall to authenticate against Facebook. This is experimental and is not suggested for deployments.
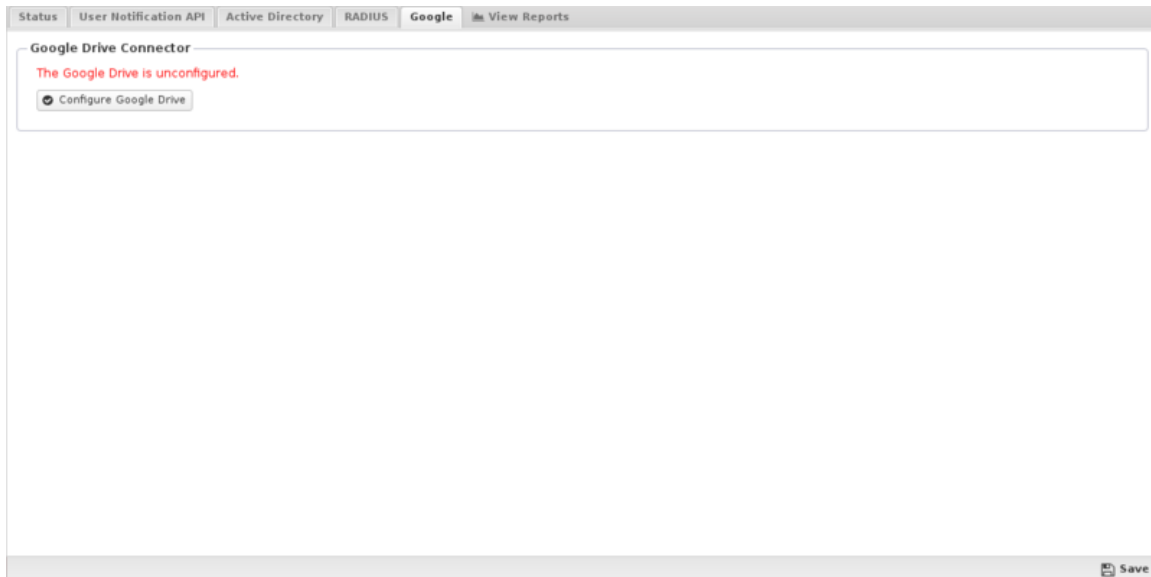
Figure 40: Google Drive Connector

Reports

The **Reports** tab provides a view of all reports and events for all sessions handled by Directory Connector.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Directory Connector Summary	A summary of Directory Connector actions.
API Usage	The amount of login, update and logout user notification API events over time.
API Events	Events from the user notification API.

The tables queried to render these reports:

- Directory Connector Reports

Related Topics

Policy Manager

Captive Portal

7.1.1 Directory Connector Reports

The **Reports** tab provides a view of all reports and events for all sessions handled by Directory Connector.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Directory Connector Summary	A summary of Directory Connector actions.
API Usage	The amount of login, update and logout user notification API events over time.
API Events	Events from the user notification API.

The tables queried to render these reports:

- Database Schema

Related Topics

Report Viewer

Reports

7.2   Reports

Reports provides users with detailed statistics of the traffic and activity on your network.

Contents

- About NG Firewall Reports
- Settings
  - Status
  - All Reports
- Report Entry
  - Data
  - Email Templates
  - Reports Users
  - Name Map
- Accessing Reports
- Report Viewer
- Conditions
  - Condition Operators
  - Conditions Example - Policy by Policy ID
  - Conditions Example - Web Filter Categories
- Related Topics

About NG Firewall Reports

These reports can be viewed online, either through the administration interface or through the separate reporting interface available to non-administrators reporting-only users.

Customizable report summaries can be sent via email, which includes basic information and a link to view the online reports if the user has access.

Reports can backup your data in multiple formats to Google Drive for long term storage.

Settings

This section reviews the different settings and configuration options available for Reports.

Status

On this tab you can click **View Reports** to open up Reports in a new browser tab.

Figure 41: Reports Status tab

All Reports

This is the full list of all currently existing reports. This includes all the default reports and any custom reports that have been added.

To edit a report, click on **View** and then click on **Settings.**

To delete a custom report, click on **View** and then click on **Delete.**

To create a new custom report, click **View** on an existing similar report, then click **Settings.** then change the name and click **Save as New Report**.

To create a report from scratch, go to **Reports** and click on **Create New** in the lower left. When creating reports from scratch, each field must be carefully chosen and tuned until the desired data is provided. This process can be time consuming and difficult. It is suggested to work with a similar report to required the desired result. Additionally you can ask for help via support or the forums and import the report if someone can craft it for you.

If creating a report from scratch, the settings and fields and their purposes are described below.

Figure 42: All Reports tab

Report Entry

A report is many settings that essentially describe how to craft a SQL query and how to display the data. Here are the fields:

Name	Value	Available	Description
Report Type	Text, Pie Graph, Time Graph, Time Graph Dynamic, Event List	The type of graph	
Title	Text	All	The report title
Category	Any existing category/application	All	The category in which the report is located
Description	Text	All	A brief description of the report
Text String	Text	Text	The text used to create the Text Report Type
Pie Group Column	Text	Pie Graph	The column to "group by" in top X charts (usually user, host, and so on)
Pie Sum Column	Text	Pie Graph	The column to sum in the top X charts (usually count, bytes, and so on)
Order By Column	Text	Pie Graph	The column to order by
Graph Style	Pie, Pie 3D, Donut, Donut 3D, Column, Column 3D	Pie Graph	The render style of the pie graph
Pie Slices Number	Integer	Pie Graph	The number of slices to display
Units	Text	Pie Graph	The units being displayed (usually bytes, sessions, and so on)
Graph Style	Line, Area, Stacked Area, Column, Overlapped Column, Stacked Columns	Time Graph	The render style of the time graph

Name	Value	Available	Description
Time Data Interval	Auto, Second, Minute, Hour, Day, Week, Month	Time Graph	The time aggregation unit or resolution
Approximation	Average, High, Low, Sum	Time Graph	The method used to aggregate/combine data points
Units	Text	Time Graph	The units being displayed (usually bytes, sessions, and so on)
Series Renderer	None, Interface, Protocol	Time Graph	The renderer used to display human-readable names
Dynamic Column	Text	Time Graph Dynamic	The column to select for/ group by
Dynamic Value	Text	Time Graph Dynamic	The value to sort by and display
Dynamic Limit	Integer	Time Graph Dynamic	The number of series to show
Aggregation Function	Count, Sum, Min, Max	Time Graph Dynamic	The function used to aggregate by dynamic values grouped by dynamic column
Graph Style	Line, Area, Stacked Area, Column, Overlapped Column, Stacked Columns	Time Graph Dynamic	The render style of the time graph
Approximation	Average, High, Low, Sum	Time Graph Dynamic	The method used to aggregate/combine data points
Units	Text	Time Graph Dynamic	The units being displayed (usually bytes, sessions, and so on)
Series Renderer	None, Interface, Protocol	Time Graph Dynamic	The renderer used to display human-readable names
Colors	Color Picker	All	The color palette to use
Display Order	Integer	All	The integer used to determine the report's position in the category list

Data

Data Retention: This value controls how much time report data is kept on disk. Please note that increasing the number increases the amount of disk space that is needed for data storage.

> **Note:** NG Firewall version *16.3* and above stops reporting data when free space falls below *5* GB.

- *Delete All Reports Data*: If you are running low on disk space and you wish to free space by wiping the reports database you can use this option.

Google Drive Backup: If your system is connected to your Google account, you can configure Reports backups to Google Drive.

- *Upload Data to Google Drive* If enabled, and the Google Connector in Directory Connector is enabled, your daily data will be uploaded to google drive each night for safe storage.
- *Upload CSVs to Google Drive* If enabled, and the Google Connector in Directory Connector is enabled, your daily CSV files will be uploaded to google drive each night for safe storage.

- *Google Drive Directory* configures which subdirectory data will be uploaded to in google drive.

Import/Restore Data Backup Files imports data from a previous backup into the database.

> **Note:** This directly imports the SQL contents. If you have upgraded and the database schema has significantly changed since the time of the back, the import will not work correctly.

Figure 43: Reports - Data tab

Email Templates

You can customize emailed reports using Report Templates. You can create as many as you wish with any combination of:

- **Interval**: Daily, Weekly, Monthly, Week to Date, Month to Date. The determines the time interval that the report will cover. Beware that enough data is available via the Retention settings to provide the data for the configured interval.
- **Mobile**: Generate chart images more appropriate for a mobile device.
- **Reports**: Select those reports under Config and Application sections. Text and chart reports are allowed but not event list reports. Reports for applications will be included only if that application is installed.

Additionally, you can copy the settings for an existing report.

The default Daily Reports template includes common text and chart reports for your system. This template is fixed and cannot be changed or modified.

Email Templates must be associated with Report Users.

Figure 44: Email Templates tab

Reports Users

Reports users are users that are not administrators, but can still view reports.

- **Email Address** is the email address (and username) of the report user. *Admin* is a special case that determines if administrators will receive emails and alerts.
- **Email Alerts** determines if this reports user will receive email alerts.
- **Email Reports** determines if this reports user will received email report summaries.
- **Email Templates** determines which email report summaries this user will receive if *Email Reports* is enabled.
- **Online Access** if enabled, a URL to online reports is included in emailed report summaries for this user.
- **Change Password** changes the password for this reports user.

Figure 45: Reports Users tab

Name Map

You can use the Name Map to manually configure the hostname for hosts. NG Firewall often can automatically determine the hostname for the IP automatically via DHCP or other methods. You can view the current names for currently active hosts in the Hosts.

However, when NG Firewall is unable to automatically determine a hostname for an IP the Name Map provides a way to manually name them.
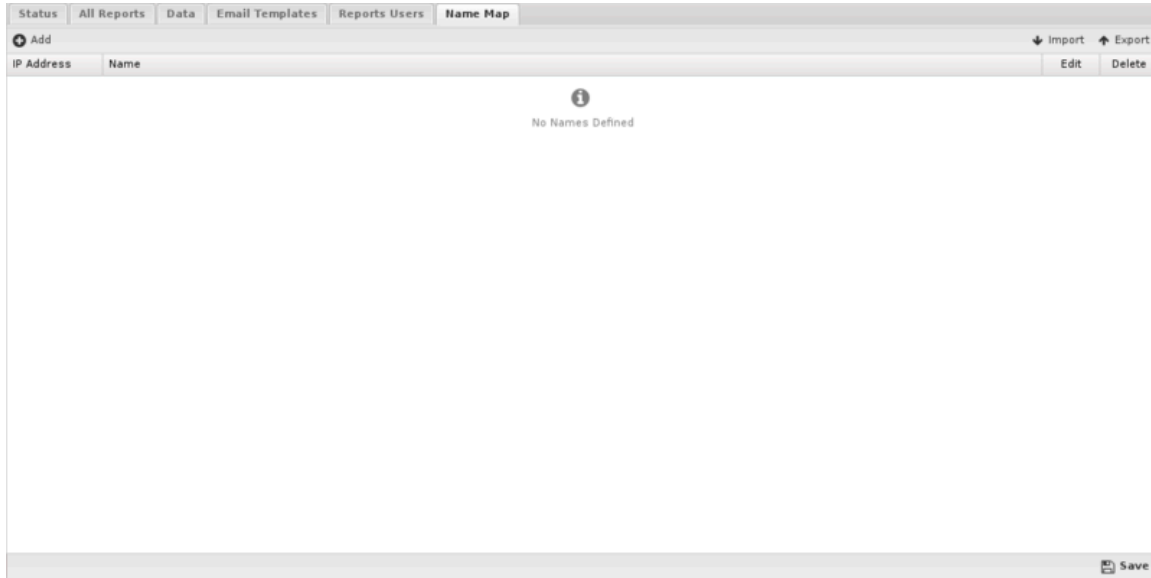
Figure 46: Name Map tab

Accessing Reports

If a user is set up to receive email report summaries, they only need to view or download the HTML attachment to see an overview report. If they need more information or would like to drill down to specific users or machines, they can use the link in the email, which will open Reports on the NG Firewall if it is accessible from their location.

To access Reports directly from a browser, you have two options:

- **Outside the NG Firewall network**: Browse to the IP of the NG Firewall /reports using HTTPs, such as https://1.2.3.4/reports.
- **Inside the NG Firewall network**: Browse to the IP of the NG Firewall /reports, such as https://192.168.1.1/reports.

Note that to view Reports from outside the network you'll need to check **Allow HTTPS on WANs** at **Config > Network > Advanced > Filter Rules**. If you have changed the **External HTTPS Port**, you'll need to use the proper HTTPS port when connecting from the outside.

Figure 47: Access Reports Summary

Report Viewer

Reports provide a graphical view of the network traffic and actions of your NG Firewall. Various reports are available within applications and base system components. The reports can be manipulated to drill down, customize, and export data in many ways using the Report Viewer.

There are a few panels in the Report Viewer:

- **The top panel**: This top panel (just below the navigation menu) allow you specify which data is viewed. By default, there is just a timeframe and no conditions, so reports will show all data for the specified timeframe. Conditions can be viewed to view more specific data, such as a specific host, user, domain, application, web category, etc.
- **The left panel**: This allows you to choose the report you wish to view. At the bottom you can use the search box to quickly find reports with the specified string in the title. You can also import and create new reports using the "Add/Import" button.
- **The chart panel**: This panel shows you the specified report. It also includes several action buttons at the top.
- **The data panel**: The data panel, hidden by default, can be displayed by clicking on the **Data View** button in the chart panel. This will show the raw data used to generate the chart and allow the user to export the data by clicking the **Export Data** button at the bottom.

Conditions

The Conditions panel appears at the top panel and can be used to filter data displayed in reports. For example, to view a "specific" host's report, you can add a condition for Client = "*192.168.1.100*" and then all reports available will only show data where the client is *192.168.1.100*. Multiple conditions can be added to drill down and inspect data. Conditions can also be added quickly by clicking on slices in pie charts.

The Add Condition dropdown contains many commonly used conditions, or the full list of all tables and columns can be browsed by clicking on the **More** button to add conditions for any database column.

Note: Conditions will not apply to all reports. For example, If viewing a specific users report by adding a condition where *Username* = *foobar* - many reports will be greyed out and unviewable. This is because the data used to generate those reports is not relevent to the specific user (it

does not contain a username column). For example, the CPU usage report is a system report that is not relevant to a specific network user and so there is no way to filter that data by user.

Condition Operators

The second field in the condition is the logical operator that will be used in evaluating the condition value defined in the last field. In most use cases the default "=" operator is what you want to use. However, there are several other operators available that make the reports and alerts a whole lot more powerful.

Conditions Example - Policy by Policy ID

In many cases, you may just want to see the traffic related to a specific policy within Policy Manager. This can be accomplished very easily by adding a condition using the Quick Add feature.

1. In the Conditions panel, select **Add**.
2. Choose **Policy ID** and specify equals and the policy ID in question.
3. The conditions is applied and will remain applied as you switch between reports.

Conditions Example - Web Filter Categories

From pie charts, you can quickly add a condition from the Current Data window. This can be handy for use with the Web Filter category selection which we'll use for this example. Once the condition is applied, we can then use other reports to drill down to find out more information about the traffic such as which user might be responsible.

1. Open Report Viewer or the Web Filter Reports tab.
2. Select the **Top Categories** report (by size or requests). In our example, you can see Games was at the top.
3. Click on the Games pie slice, and when prompted to add a condition click **Yes**.
4. All Reports can now be viewed for Games only traffic.
5. For example, the Top Clients (by request) will show the clients that visited the most gaming sites.
6. For example, the Web Usage (scanned) will show "Gaming" web usage throughout the day of the network.

Related Topics

Custom Reports

7.2.1   Custom Reports

Custom Reports allow you to create a report to your specifications and save the report for future use. These can be added by copying an existing report and modifying it to your needs, or by creating a report from scratch. Both methods require a strong understanding of how reports in Arista work and the Database Schema.

If you're having trouble creating any reports give us a shout on the forums or through support. As we receive requests for common reports we will add these as default reports in future releases.

Add Custom Reports

The Copy or Customize features allow you to use an existing report as a base for creating new reports. We recommend using an existing report to base your custom report as the fields and queries are much easier to understand.

Copy Report

1. Go to the **Reports** Section,
2. Locate the report you want to copy and click **Settings**.

3. Rename the title and click the **Create New** button. A copy of the report is made with the new title.
4. Modify other report settings as needed.
5. Click **Preview/Refresh** to view changes.
6. Click **Update** to save the changes to the report.

Customize Report

1. Rename the report and modify other report settings as needed.
2. Click **Create New** to save the report with a new name.

Add Report

1. Go the **Reports** section.
2. Click the **Add/Import** button and select **Create New**. A blank report form will appear.
3. Name the report and description, select the category and report type of graph and fill other fields to customize the report.
4. Click **Preview/Refresh** at any time to preview the report.
5. Click **Create New** or **Update** to save the report.

Report Fields
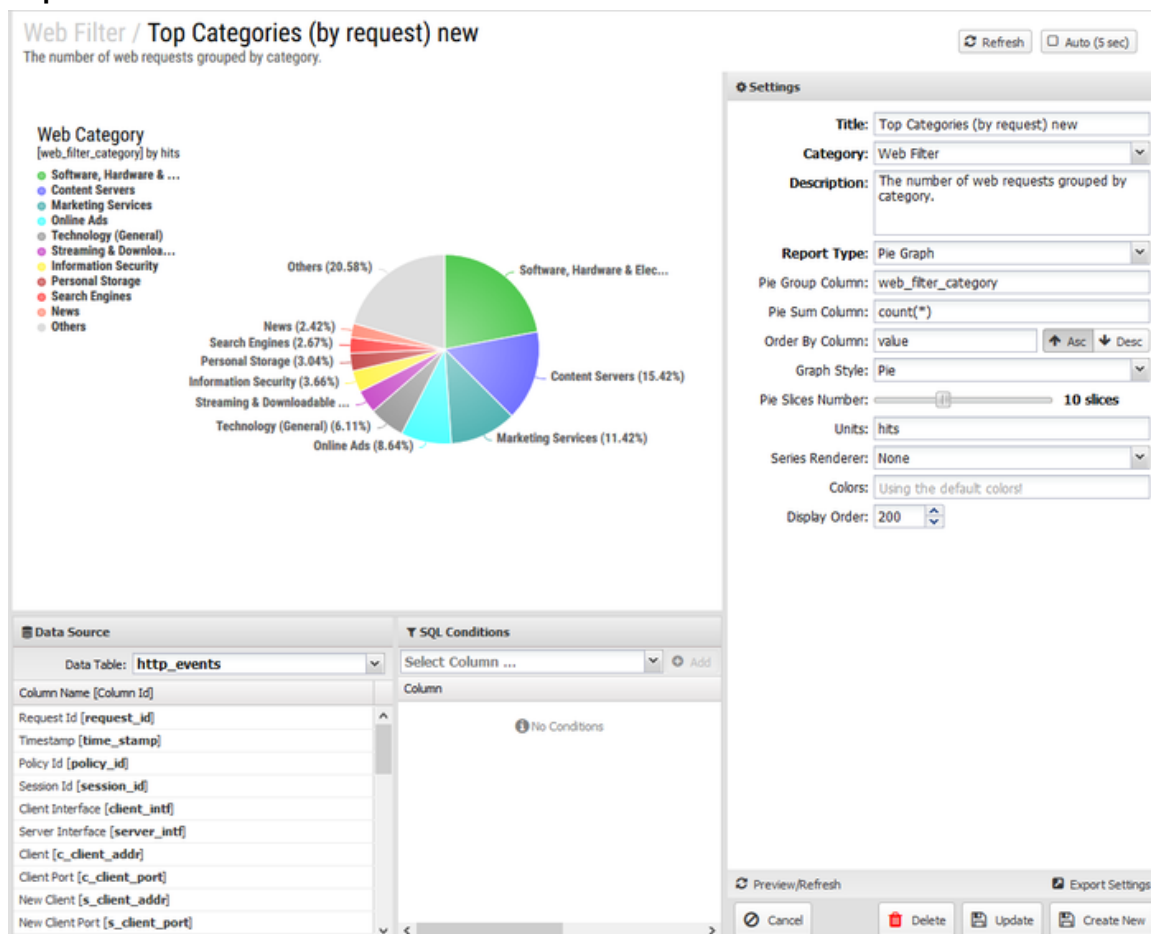
Figure 48: Custom Report Fields

The following fields will be there when modifying or creating a new custom report.

Report Field	Description
Category	The app where the report will be displayed.

Report Field	Description
Title	Title of the report.
Description	Description will be shown under the title of the report to provide the reader more information.
Enabled	Disabled reports will not be shown in the list within the application.
Display Order	Determines the order the reports are displayed within the application.
Units	Unit of measure displayed on the graph.
Table	Table from the Database Schema to query.
Type	Time Chart (Line or Bar), Pie Chart, or Text Chart.
Time Chart Style	Line, Bar, or Bar 3D. Bar charts can also be overlapped or separate columns.
Time Data Interval	Auto is recommended. You can also use seconds, minutes, hours, day, week or month.
Time Data Columns	Data from the table defined above that you wish to use for your chart. This is written with SQL syntax. Add each column or data series on a new line. See the Web Usage (all) report for a good example.
Pie Group Column	Column from the table defined above that you wish to group data by. This uses the column name from the Database Schema.
Pie Sum Column	Determines how data is calculated after grouping. Common functions would be count() and sum() for specific columns within the table. See the Web Filter Top Sites reports (by size and by request) for good examples.
Pie Slices Number	Determines how many results are shown.
Text Columns	Data from the table defined above that you wish to use for your page. This is written with SQL syntax. Add each column or data series on a new line. These will be referred to in the text string using the line number.
Text String	Text string to be shown on the page. This can include columns from the text columns query. See the Ad Blocker summary report for a good example.
Color	Colors used in the chart, defined by hex color codes.
Order By Column	Determines how the order will be determined. This will typically be 'value' for pie charts, and left blank for others.
Order Direction	Order for the data to be displayed. For a pie chart, this is typically Descending and for time chart, this is typically Ascending.
Sql Conditions	Additional conditions added to the data. This can be useful for specifying a username, IP address of other information.

Custom Report Example - User Reporting

A commonly requested report is to be able to monitor a problem employee or student who is using bandwidth excessively or accessing inappropriate web content. You can add this report so that it is always accessible under the Web Filter report tab very easily.

To start, get an idea of the reports you're interested in. For this example we'll use Top Sites (by size), but others like Top Categories or Top Flagged sites might also be useful. This is intended to be a rough outline of the procedure; adjust as needed.

Creating the Report

1. Go to **Reports**
2. Scroll to the Web Filter and select a report.
3. Update the Title for the report and description if necessary.
4. Click **Create New**
5. Since you are interested in the same type of data, all other report fields can remain the same.
6. In SQL Conditions, add **conditions**. For this example we will use the **username** but you might also use hostname, client address of other fields.

    a. In the **Select Column** field, select **username**.
    b. Click **Add**.
    c. Use the = operator and enter the username for the user in the **Value** field.
7. Click **Update**.
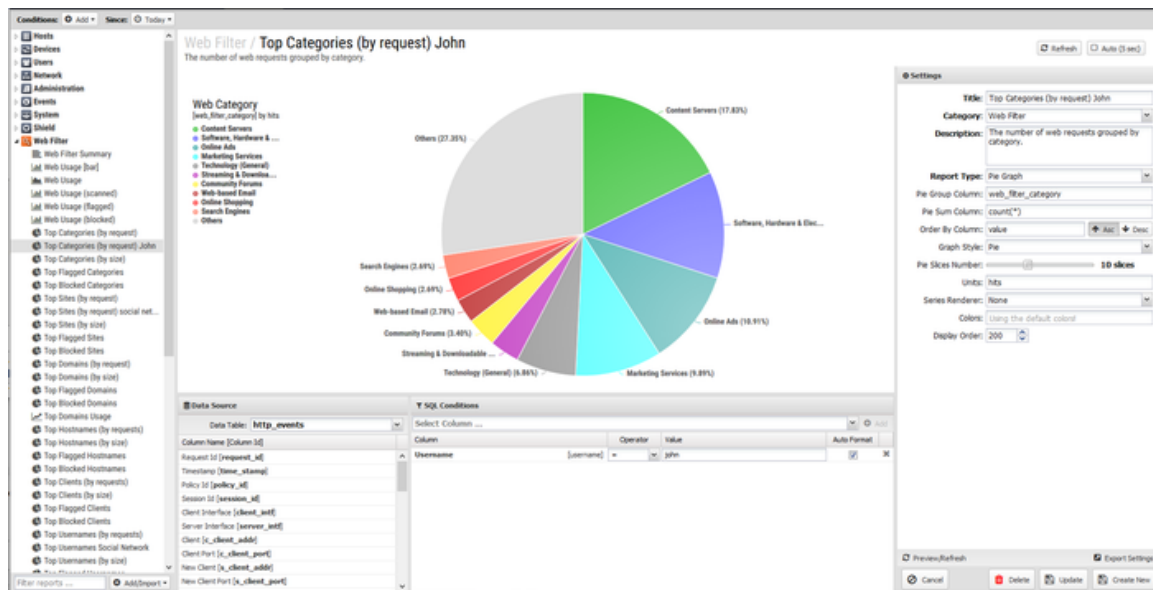
Figure 49: Custom Report

Custom Report Example - Rack Reporting

With Policy Manager, racks can be created to allow different policies to apply to different groups. When reporting, you might wish to see how traffic is being handled differently across racks. Creating custom reports to show information related to a specific rack is easy to accomplish.

For this example we'll use Application Control Top Flagged Applications, but this same procedure would apply for any other application or report. This is intended to be a rough outline of the procedure; adjust as needed.

Find the Rack ID

1. Open **Policy Manager** Settings.

2. Under the Policies tab is a listing of racks with the ID.
3. Take note of the rack(s) of interest. For this example we'll use rack *id 6* for Mobile Devices.

Figure 50: Rack IDs

Creating the Report

1. Go to **Reports**.
2. Scroll to the **Application Control** reports and select a report.
3. Update the Title for the report and description. Here we have added "Mobile Devices Rack".
4. Click **Create New**
5. Since you are interested in the same type of data, all other report fields can remain the same.
6. In Sql Conditions, add the condition for the policy_ID.

   a. Click **Add**.
   b. In the **Column** field, select **policy_id.**
   c. Use the = operator and enter the policy id of 6 (matching the Mobile Devices rack) in the **Value** field.
7. Click **Done**.
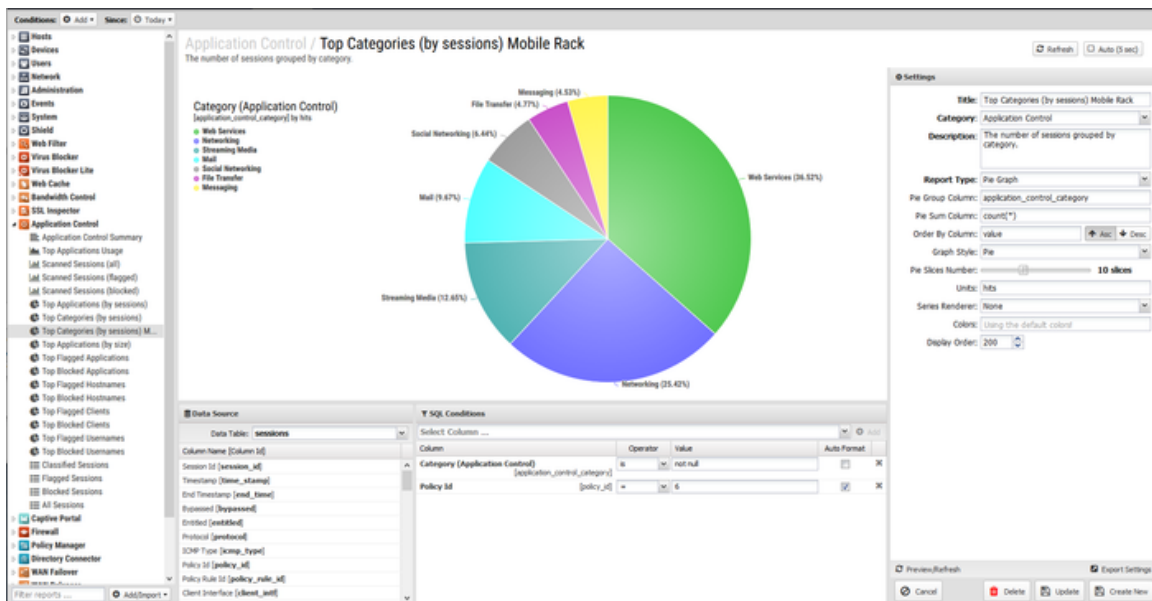
Figure 51: Custom Report

7.3   Policy Manager

Policy Manager is one of NG Firewall's most powerful features. It works by allowing you to create multiple policies, which are different sets of applications configured differently for different use cases.

About Policy Manager

Often you will configure all of NG Firewall apps to service the whole network, however often it is necessary to handle traffic for some users or network devices differently. For example, you may want Web Filter to be different for students vs teacher and you may want no Web Filtering at all for your servers. You may wish to run Captive Portal but only on the WiFi network or for unidentified devices. You may wish block certain applications with Application Control, but only at certain times of days.

In these cases, Policy Manager can simplify configuration a great deal by allowing for multiple sets of applications to be configured differently. Policy Manager allows creation of new policies beyond the "Default Policy". To address the above examples, the administrator can create a "Student Policy", "Teacher Policy", "After hours Policy", and "Wireless Network Policy" etc. Each policy can run a different set of applications configured differently. The Policy Manager Rules can determine which policy handles which network sessions.

- Set up multiple policies for different users, hosts, networks, interfaces, times of day, days of week, etc.
- Choose what applications are running in each policies.
- Configure multiple applications in separate policies simultaneously using the **Parent Policy system.**

  This allows you to "copy" the configuration of *some* applications from another policy, but not others - this makes doing things such as having different Web Filter settings across policies, but keeping the configuration of all other applications identical across policies. There is not usually a need to modify settings for applications like Virus Blocker or Spam Blocker between different user groups, however if it is necessary it only takes a few clicks.

Note that we will be using the example of a school a lot in this section as it is quite apt in showing how Policy Manager can help you with different user classes. This can be applied to any organization; just look for groups you can fit users into - Administrative Assistants, Marketing, or HR, you're free to choose. It can also apply to different sets of servers (*i.e.*, a *DMZ policy* for handling public servers, and an *Internal policy* for handling internal user machines and a *Wireless policy* for handling wireless users) and it can also apply to different times of day (*i.e.*, a *Lunchtime and After Hours policy* and a *Work Hours policy*). For simplicity the examples below will mostly use the school groups as an example.

Getting Started with Policy Manager

Policies provide a way to handle different settings for different sessions. Using our example, an NG Firewall protecting a school might have three different policies - Students, Teachers, and Administrators. These policies provide completely separate configurations for traffic processing, for example you could allow teachers to access Facebook but not students.

NG Firewall will always have at least one policy, the **Default Policy**. You can rename but you cannot remove this policy. As mentioned before, you create Rules to send traffic to **policies** where it is processed by the **applications**. Policies are created from within Policy Manager, however you will use NG Firewall's web GUI itself to switch to and configure each policy. At the top of the web GUI, you will see **Default Policy** with an arrow next to it - clicking this arrow allows you to change the policy you're looking at, as well as access the Sessions Viewer and Hosts Viewer and open the Policy Manager Settings directly.

Parent Policies

When you first create a new policy it will not contain any applications. You can add in any applications you want and configure them to your liking, or you can use the **Parent Policy** system. When creating a new policy using Policy Manager, you have the option to select a Parent Policy. If you use this option your new policy will be pre-populated with all applications and settings from the Parent Policy you selected, however it will look a *bit* different.

When you view the new **child** policy, the application faceplates will be greyed out and you will be unable to click Settings. This is because the settings for these applications are being *inherited* from the parent policy, which is useful because it saves you from having to reconfigure applications you want operating the same way in multiple policies, such as the virus scanners. To change the settings or view the Event Logs, you'll need to open the application on the parent policy and use the drop-down to select the policy to view traffic for.

If you want to modify the settings of an application in a child policy, you'll need to install the application you want to modify in the child policy - I know, it's already there, but you can't click Settings to modify the configuration. On the **Apps** tab on the left side of the web GUI, just click **Install** again. After a few seconds the app will re-appear and you will be able to click settings. Once this has happened, the new child application *overrides* the application inherited from the parent. The settings of the parent policy for this application **have no effect for the application you have added to the child policy**. If you're following along, your child policy will contain all applications that your parent policy does, however only one will have the Settings button enabled.

To recap using our school example, we would send students to the Default Policy, then create a new Teacher Policy which uses the Default Policy as its Parent Policy. If you go to the Teacher Policy, all the apps will be greyed out and you will not be able to modify any settings because they are copied from the Default Policy. By adding in a new copy of Web Filter, you can modify the settings so the teachers can access websites the students cannot, however settings for all other applications will still be copied from the Default Policy.

Settings

This section reviews the different settings and configuration options available for Policy Manager.

Status

This displays the current status and some statistics.
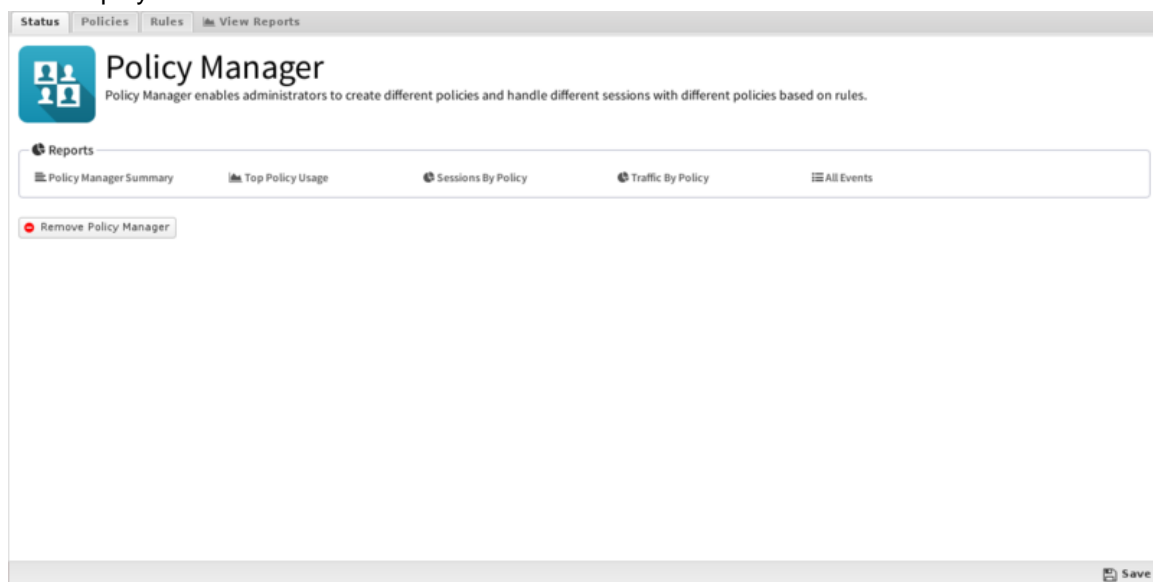
Figure 52: Policy Manager Status

Policies

From this tab you can create new policies, however please note you'll first need to create and save a policy before you can create a rule to apply traffic to that policy.

To create a new policy, simply click Add in the **Policies** section.

- **Name**: The name of this policy as displayed in the web GUI.
- **Description**: A description for this policy.
- **Parent Policy**: Which policy (if any) this one should use as a Parent Policies.
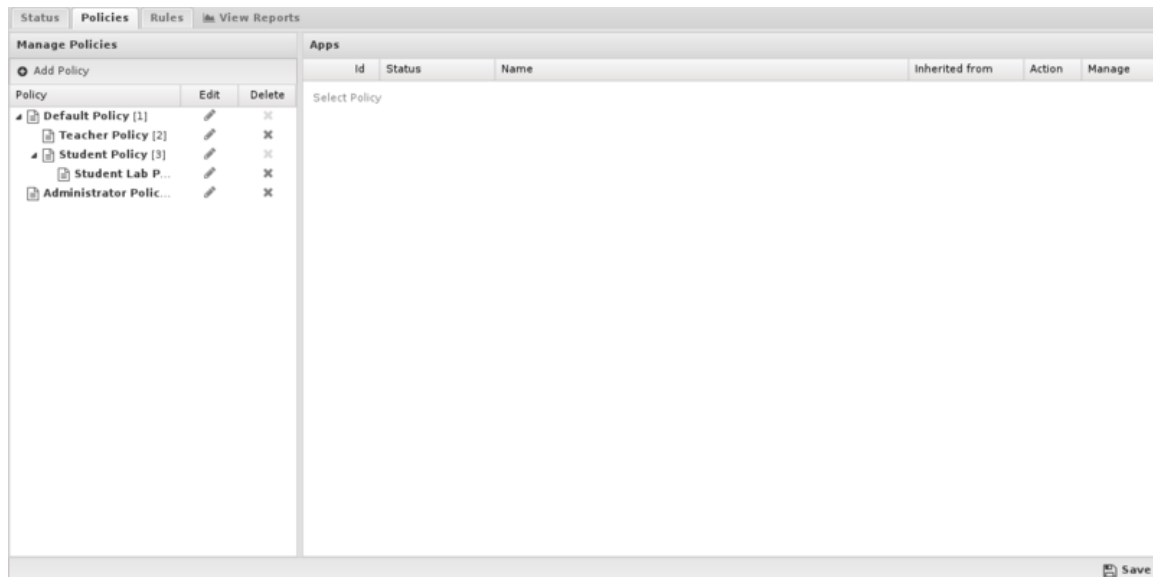
Figure 53: Policy Manager Policies

Rules

If you've been reading up until this point, you may have guessed that this new policy will do nothing until you send traffic to it. To accomplish this, you'll need to create a rule - click Add in the **Rules** tab.

When each new session is processed, the rules are evaluated in order. If all of a sessions attributes match all of the criteria of a rule it is considered a match. The policy for the first matching rule will be used to process the session. If no rules match, the *Default Policy* will be used to process the session.

These rules operate as described in the Rules documentation.

Like many areas of NG Firewall, the rules work from the top down. Let's go back to our school example and say we have three policies: Default (for students), Teacher and Administrative Staff. To get traffic to these policies, we would need to create two policies: one for the Teacher policy and one for the Administrative Staff policy - any traffic that did not match those two policies would be sent to the Default Policy. You can also explicitly add a rule sending traffic to the Default Policy, although it's not required.

If the policy rule for the Teacher policy is incorrect, it may end up matching *all* network traffic and sending it to the Teacher policy. Because it matches, the rule under it (for the Administrative Policy) will never be evaluated. On the flip side, if a rule is too narrow, it may not match the traffic you're trying to match at all, dumping it on the Default Policy. Because of this, we recommend starting out policy rules as very general (*e.g.*, match a few IPs or an entire subnet) and then tightening them down from there - let's look at some common example policies:

File:Policy rules.png Quick look at the rules table

So this is the **Rules** tab of the Policy Manager settings. On this page you will be able to see all of the rules that will be evaluated when determining which policy to direct a specific user or IP address. Pay

attention to the 'Target Policy' column. Any users defined in the corresponding rule will be directed to that policy. Click on the page icon under the 'edit' column to define the rule.

File:Policy rules2.png Click **Edit** and define the rule.

Once you've clicked the edit button, this is where you'll end up. This is where you will define which users will be assigned to which policy. You can specify a user using any one or combination of the identifiers in the drop down box. Any users that match the specified identifiers will be directed to the policy specified in the `Target Policy` field.
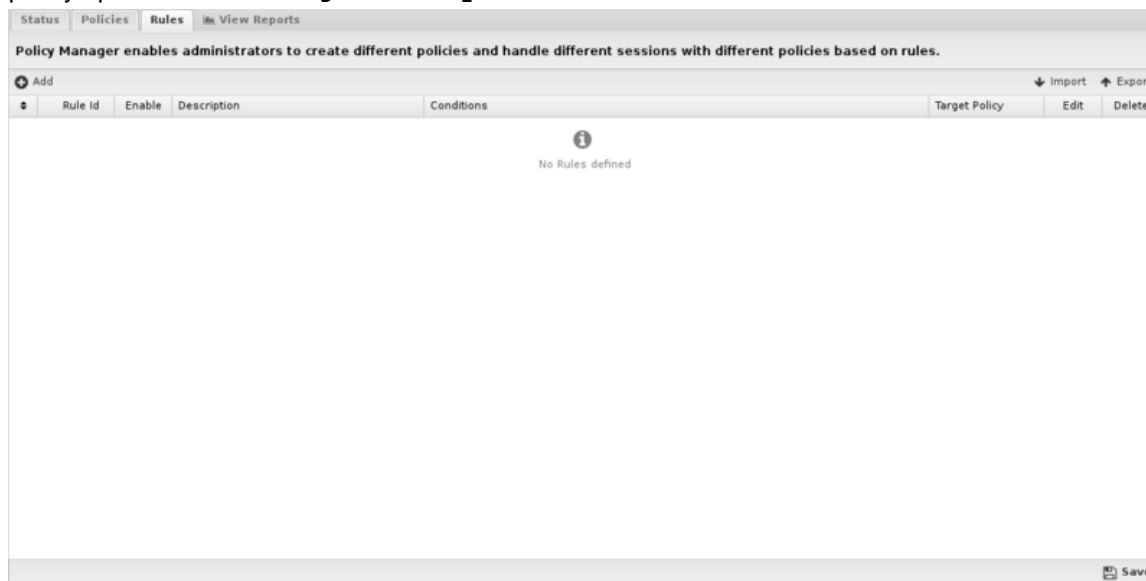
Figure 54: Policy Manager Rules

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Policy Manager.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Policy Manager Summary	A summary of Policy Manager actions.
Top Policy Usage	The amount of bandwidth per policy.
Sessions By Policy	The number of sessions for each policy.
Traffic By Policy	The amount of traffic for each policy.
All Events	Lists all sessions with the Policy Manager policy that handled the session.

The tables queried to render these reports:

- Database Schema

Related Topics

- NG Firewall Rule Syntax

- Directory Connector

7.3.1 Policy Manager Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Policy Manager.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Policy Manager Summary	A summary of Policy Manager actions.
Top Policy Usage	The amount of bandwidth per policy.
Sessions By Policy	The number of sessions for each policy.
Traffic By Policy	The amount of traffic for each policy.
All Events	Lists all sessions with the Policy Manager policy that handled the session.

The tables queried to render these reports:

- Database Schema

Related Topics

Report Viewer

Reports

NG Firewall Filter Apps

Contents

- Ad Blocker
- Application Control
- Application Control Lite
- SSL Inspector
- Spam Blocker
- Spam Blocker Lite
- Web Filter
- Web Monitor

8.1 Ad Blocker

Ad Blocker allows you to block a majority of advertising content that is delivered to users on web pages that they request.

Ad Blocker uses downloadable filter subscriptions from a variety of sources which contain lists of web sites and extensions that are typically used to deliver advertising.

⚠️ **Warning:** Ad Blocker is similar in nature to browser plugins that block ads (like AdBlockPlus, Ghostery). Like those plugins, sometimes blocking ads and tracking can interfere with the proper functioning of online websites or media. However, Ad Blocker blocks at the gateway the browser user has no visual indication that something has been blocked nor have any way to temporarily disable or bypass the blocking if it interferes with the proper functioning of online services. As such, Ad Blocker can cause issues and the browser user has no recourse. Running Ad Blocker can impose administrative overhead dealing with issues and is ideal for some environments (small sites, home, enthusiasts, etc) but not recommended in others like schools, businesses, and large sites.

⚠️ **Warning:** With the increasing adoption of SSL ad blocker can do very little without SSL inspection, but running SSL inspection is not ideal for many organizations.

⚠️ **Warning:** When the update button is pressed it pulls new signatures directly from 3rd party sources. We have not tested any future updates and can not guarantee any future updates will work correctly. Updates may interfere with the proper functioning of websites and potentially cause massive problems.

⚠️ **Warning:** Unlike other apps, Ad Blocker is off by default after installation. Do not enable Ad Blocker until reading the above warnings.

Settings

This section reviews the different settings and configuration options available for Ad Blocker.

Status

This displays the current status and some statistics.
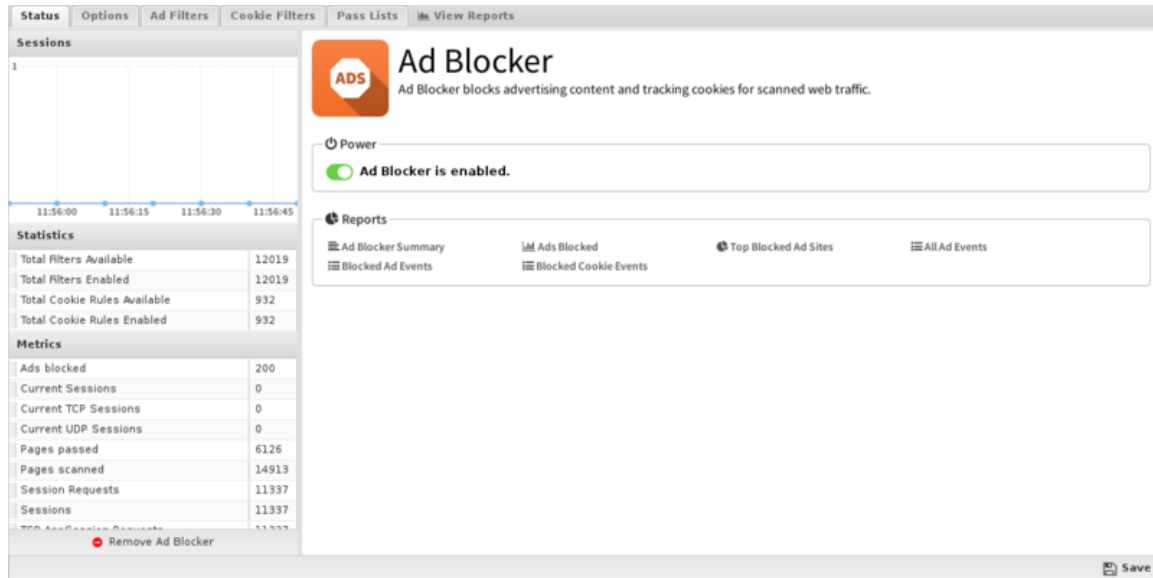
Options

If **Block Ads** is enabled then Ad Blocker will block web requests that it determines are for advertisements. It will return a **HTTP 403** permission denied to the client that requests the ad.

If **Block Tracking & Ad Cookies** is enabled Ad Blocker will block cookies that are used for advertising or behavior tracking purposes.

The **Update** button will update the ad filters signatures.

> ⚠ **Warning:** When the update button is pressed it pulls new signatures directly from 3rd party sources. We have not tested any future updates and can not guarantee any future updates will work correctly. Updates may interfere with the proper functioning of websites and potentially cause massive problems.

Ad Filters

Ad Blocker's **Standard Filters** list will populate with many entries to match common ad serving strings. This list can not be modified, but you can enable or disable the rules here. You are free to add to and edit rules in the **User Defined Filters** tab, as you see fit - simply click **Add** and enter a description to match. Be careful when selecting blocking criteria as you may block much more content than planned if your criteria is not carefully specified.
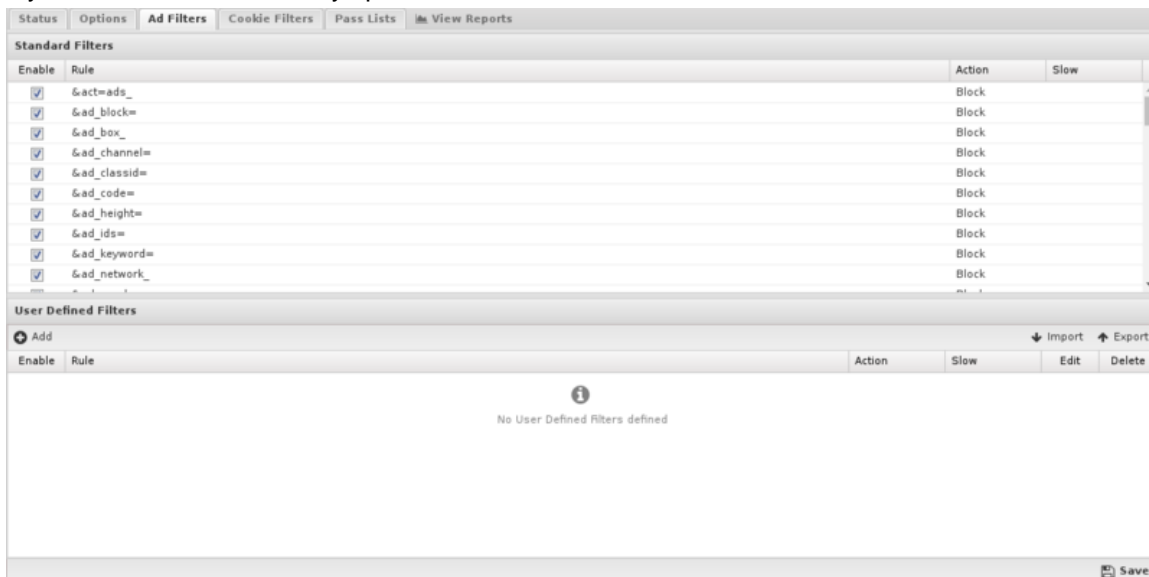
Cookie Filters

The **Standard Cookie Filters** list is populated with entries to match common cookie domains. This list can not be modified but you can enable or disable the rules here. You are free to add to and edit rules in the **User Defined Cookie Filters** tab, as you see fit.
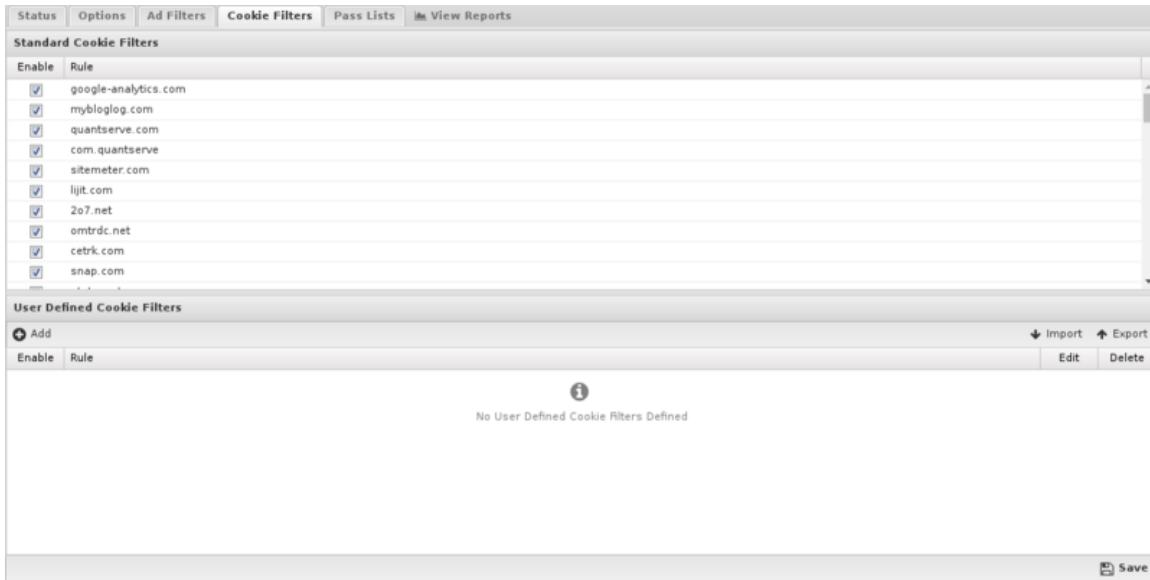
Pass Lists

Pass Lists are used to pass content that would have otherwise been blocked. This can be useful for "unblocking" sites that you don't want blocked or allowing certain users special privileges.

- **Passed Sites**: Any domains you add to the Passed Sites list will allow ads and cookies, even if blocked by an existing filter - just add the domain and save. Unchecking the pass option will allow ads to be blocked as if the entry was not present.
- **Passed Client IPs**: If you add an IP to this list, Ad Blocker will not block any ads or cookies from that IP. Just add the IP and save. Unchecking the pass option will have the block/pass lists affect the user as if they were not entered into the Passed Client IPs list.
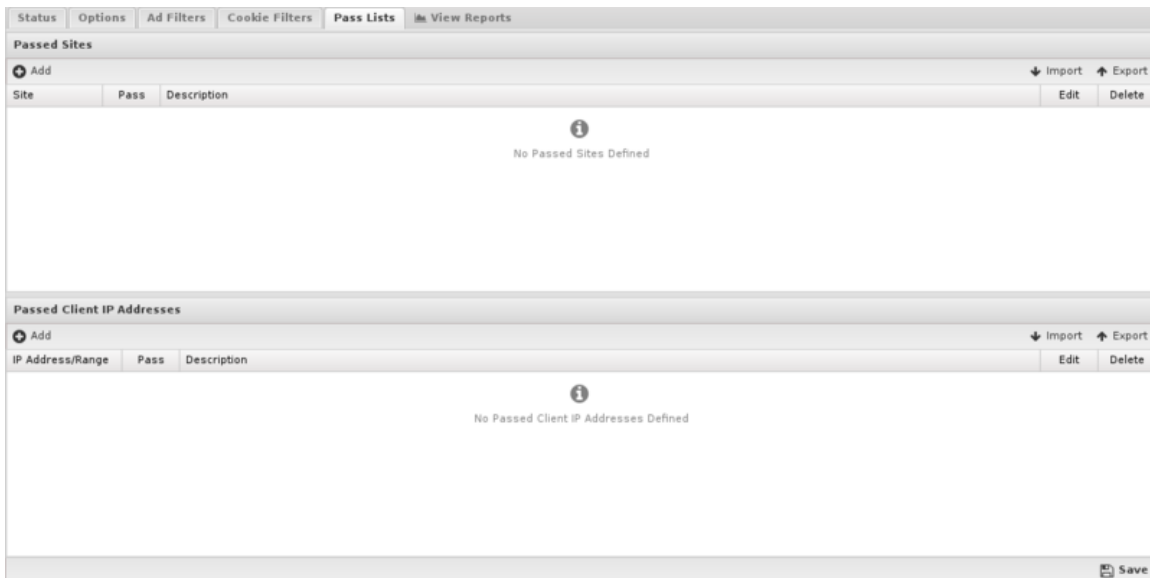
Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Ad Blocker.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries: report Entry	Description
Ad Blocker Summary	A summary of ad blocker actions.
Ads Blocked	The amount of detected and blocked ads over time.
Top Blocked Ad Sites	The number of blocked ads grouped by website.
All Ad Events	All HTTP requests scanned by Ad Blocker.
Blocked Ad Events	HTTP requests blocked by Ad Blocker.
Blocked Cookie Events	Requests blocked by cookie filters.

The tables queried to render these reports:

- Database Schema

Related Topics

- Report Viewer
- Reports
- Web Filter
- Phish Blocker

8.1.1   Ad Blocker Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Ad Blocker.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Ad Blocker Summary	A summary of ad blocker actions.
Ads Blocked	The amount of detected and blocked ads over time.
Top Blocked Ad Sites	The number of blocked ads grouped by website.
All Ad Events	All HTTP requests scanned by Ad Blocker.
Blocked Ad Events	HTTP requests blocked by Ad Blocker.
Blocked Cookie Events	Requests blocked by cookie filters.

The tables queried to render these reports:

- Events

Related Topics

8.2   Application Control

Application Control leverages the Network Application Visibility Library (NAVL) from Procera Networks [1] to perform deep packet (DPI) and deep flow (DFI) inspection of network traffic. This allows the server to accurately identify thousands of today's common applications such as Social Networking, P2P, Instant Messaging, Video Streaming, File Sharing, Enterprise Applications, Web 2.0 and much more.

About Application Control

For most common applications you can simply go the list on the Applications tab, check **Block** for anything you want to stop, then Application Control will take care of the rest. If you need a higher degree of control you can use the Rules tab to create custom rules which target more complex traffic patterns.

How It Works

Application Control feeds each chunk of data to a classification engine as it passes through the application. The classification engine continues to analyze the traffic flow and keeps properties of the session, such as the Application property. Each time the classification of the Application property is updated, the Applications settings are checked to see if that application is allowed. If the application is configured to be blocked in the settings the data is blocked. If not, the process continues until the session reaches a fully classified state where the classification engine believes no more classification of the session is possible. At this point the Rules are evaluated and the session is ultimately blocked or passed based on the rules you've configured.

Settings

This section describes the different settings and configuration options available for Application Control.

Status

The Status tab displays a summary of traffic and configuration information. The **Traffic Statistics** section displays the total number of sessions that have been scanned, along with the number of those sessions that were allowed, flagged, or blocked. The **Application Statistics** section shows you the total number of applications that can be detected by the application, along with the number of those protocols that will be flagged and/or blocked. **Rules Statistics** allows you to quickly see how many custom rules you have configured, as well as how many of those rules are active.
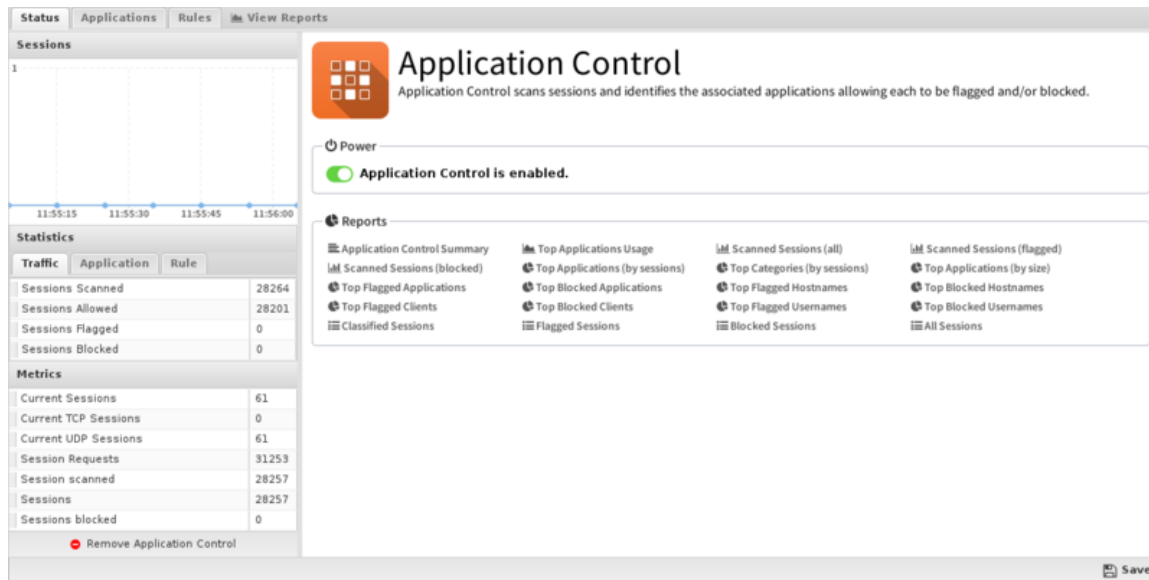
Figure 55: Application Control Status

Applications

The **Applications** tab is the primary and preferred way for using Application Control to manage network traffic. Simply find the application you want to target, and use the block and flag checkboxes as appropriate. You can sort the list on any of the columns displayed, which should help in finding and managing the protocols you want to target. Simply check **Block** to stop these applications or **Flag** to allow them, silently filing them as violations in the Reports. Use the following definitions to set up the Applications tab for your organization:

- **Application**: The unique identifier for the application.
- **Block**: Enable this checkbox to block/reset sessions of this application. For TCP this will actively reset the connection. For UDP it will drop the packet and kill the session.
- **Tarpit**: Enable this checkbox to block/tarpit sessions of this application. For TCP, this makes it appear to both the client and the server that the other party is receiving the data, but it is not responsive. It silently drops the data. For **UDP**, it is identical in behavior to block except the connection is kept open so the next packet will be dropped instead of recategorized as a new session.
- **Flag**: Enable the checkbox to flag the traffic. It will be flagged as a violation in Reports.
- **Name**: The standard name for the application.
- **Category**: A fairly general and high level category for the application.
- **Productivity**: Productivity is best thought of as an index value between 1 and 5 that rates the potential for each application to improve or increase the overall productivity of your network users, assuming of course that listening to music and playing online games is not in their job description. So, applications with a low Productivity index (e.g. MySpace, Hulu, Zynga Games) can be expected to have a negative impact on productivity. Items with a high value (e.g. Active Directory, Network File System) can generally be viewed as critical for maintaining or improving productivity.
- **Risk**: Risk is another index value between 1 and 5 that rates the potential for each protocol or application to allow really nasty stuff onto your network. The higher the risk index, the greater the chance of letting in something that could be dangerous or destructive. So low risk items (e.g. Active Directory, Oracle, LDAP) are generally no cause for concern, while applications rated with a high risk (e.g. BitTorrent, Pando, Usenet) increase the possibility you'll find yourself spending long nights deleting pirated software and cleaning up viruses and other exploits that find their way into your infrastructure.

- **Description**: Provides a more detailed description for each application in the list. In some cases the description is much larger than will fit within the grid column, so you can click on any description to see a pop-up window with the full text displayed.

Figure 56: Application Control Applications

Rules

If the traffic you need to manage can't be handled via the Applications tab you can create custom rules that will allow you to analyze and control traffic based on much more complex patterns and conditions. For each session, the rules are **only evaluated once after the classification engine has completed analysis of the traffic**. The rules are then evaluated in order until the first match is found, at which point the configured action will be performed. If there are no matches the session will be tagged as allowed, the traffic will flow unimpeded, and no further analysis of that traffic will occur.

> ⚠️ **Important:** These rules are evaluated **when the classification engine has completed all analysis.** This usually occurs after a few packets have passed. This means the rules are useful because enough has been learned about the session that is not known at the session creation time to have powerful rules, such as HTTP information or protocol/application information. If the full classification is not completed after 15 chunks of data, then the rules are evaluated given the current information.
>
> If an application is blocked or tarpitted in the Applications tab it will be blocked immediately when it is identified before the engine has completed analysis. In this case, the rules will have **NO EFFECT** because the sessions is blocked before the rules are evaluated.

Application Control Rules provide a very powerful feature that can be used to control application usage. However, understanding how and when the rules are evaluated is essential in their usage.

Figure 57: Application Control Rules

Anatomy of a Rule

An Application Control Rule is a standard rule as documented in the Rules documentation. We'll use one of the default rule entries for Ultrasurf to help explain how Rules work. This is exactly the kind of traffic that the Rules engine was created to seek and destroy. For this particular rule, the objective is to block all traffic that: **a)** uses **_port 443_**, **b)** looks like valid HTTPS traffic, and **c)** doesn't use a valid SSL certificate. To accomplish this, we created four matchers:

1. The first matcher makes sure the rule only looks at TCP traffic.
2. The second causes the rule to only look at traffic with a destination port of **_443_**.
3. The third matcher is where the real magic starts. In this case, we created Glob matcher that looks for the /SSL tag anywhere in the Application Control/ProtoChain. (Don't worry, we'll cover globs and chains below!)
4. The fourth matcher is the frosting on the cake. We tell the rule to look at the Application Control/ Detail parameter. This is where the server name from the SSL certificate will be located when an SSL encrypted session is detected. In this case we left the Value field empty, since we're looking for cases where there is no valid certificate.

Application Detail

The Detail field will contain different types of [#Is there a list of session properties? | information] depending on the protocols detected during session classification. For matcher conditions other than those listed below, the Detail field will be empty.

Matcher	Detail Contents	Example
Application: FBOOKAPP	The name of the Facebook Application that is being accessed.	wordswithfriends
Application: HTTP	The contents of the Content-Type header in the session data coming from the server.	image/jpg
ProtoChain: */SSL*	The server name extracted from the SSL certificate used to encrypt the session.	www.gmail.com

Actions

- **Allow**: Allow the traffic.
- **Block**: When this option is selected, traffic in both directions will be silently dropped, but the session will remain active.

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Application Control.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Application Control Summary	A summary of Application Control actions.
Top Applications Usage	The amount of bandwidth per top application.
Scanned Sessions (all)	The amount of scanned, flagged, and blocked sessions over time.
Scanned Sessions (flagged)	The amount of flagged, and blocked sessions over time.
Scanned Sessions (blocked)	The amount of flagged, and blocked sessions over time.
Top Categories (by sessions)	The number of sessions grouped by category.
Top Applications (by sessions)	The number of sessions grouped by application.
Top Applications (by size)	The number of bytes grouped by application.
Top Flagged Applications	The number of flagged sessions grouped by application.
Top Blocked Applications	The number of blocked sessions grouped by application.
Top Flagged Hostnames	The number of flagged sessions grouped by hostname.
Top Blocked Hostnames	The number of blocked sessions grouped by hostname.
Top Flagged Clients	The number of flagged sessions grouped by client.

Report Entry	Description
Top Blocked Clients	The number of blocked sessions grouped by client.
Top Flagged Usernames	The number of flagged sessions grouped by username.
Top Blocked Usernames	The number of blocked sessions grouped by username.
Classified Sessions	All sessions matching an application control signature.
Flagged Sessions	All sessions matching an application control signature and flagged.
Blocked Sessions	All sessions matching an application control signature and blocked.
All Sessions	All sessions scanned by Application Control.

The tables queried to render following reports:

- Sessions
- Session Minutes

Related Topics

Application Control

8.2.1 Application Control Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Application Control.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Application Control Summary	A summary of Application Control actions.
Top Applications Usage	The amount of bandwidth per top application.
Scanned Sessions (all)	The amount of scanned, flagged, and blocked sessions over time.
Scanned Sessions (flagged)	The amount of flagged, and blocked sessions over time.
Scanned Sessions (blocked)	The amount of flagged, and blocked sessions over time.
Top Categories (by sessions)	The number of sessions grouped by category.
Top Applications (by sessions)	The number of sessions grouped by application.
Top Applications (by size)	The number of bytes grouped by application.

Report Entry	Description
Top Flagged Applications	The number of flagged sessions grouped by application.
Top Blocked Applications	The number of blocked sessions grouped by application.
Top Flagged Hostnames	The number of flagged sessions grouped by hostname.
Top Blocked Hostnames	The number of blocked sessions grouped by hostname.
Top Flagged Clients	The number of flagged sessions grouped by client.
Top Blocked Clients	The number of blocked sessions grouped by client.
Top Flagged Usernames	The number of flagged sessions grouped by username.
Top Blocked Usernames	The number of blocked sessions grouped by username.
Classified Sessions	All sessions matching an application control signature.
Flagged Sessions	All sessions matching an application control signature and flagged.
Blocked Sessions	All sessions matching an application control signature and blocked.
All Sessions	All sessions scanned by Application Control.

The tables queried to render these reports:

- Sessions
- Session Minutes

8.3   Application Control Lite

Application Control Lite scans sessions for the specified signatures and will log and/or block sessions based on their content. Many applications can be hard to block based on port as modern applications will find and use open ports automatically. Application Control Lite provides a basic functionality to block sessions based on the content (data) in the session.

Settings

This section reviews the different settings and configuration options available for Application Control Lite.

Status

The **Status** tab will show you current information on signatures Available, Logged and Blocked.

Signatures

The **Signatures** tab shows the the list of current signatures. Signatures are regular expressions that are written to match known protocols as accurately as possible. New custom signatures can be designed to match certain applications or sessions, or signatures can often be found for many existing protocols L7-filter Supported Protocols.

As the early data in each session goes from the server to the client and the client to the server, it is stored in a buffer. As each chunk of data arrives, the data is evaluated against any enabled signatures. If the signature is checked "log" then the session will be tagged and logged as having matched the specified signature. If the signature is checked "block" it will be logged and the session will immediately be closed.

Writing custom signatures can be dangerous and difficult. Usually one of several outcomes will happen when writing a block signature:

- It will not match anything. In this case, the signature needs to be fixed.
- It will block the desired protocol/application and nothing else. This is ideal.
- It will only partially block the protocol. Many multi-session protocols only have some sessions identified. This can have varying effects depending on the application.
- It will block the protocol and block other things too (false positives). This can cause major problems with the network.
- It will block the protocol and the application will adapt and use an alternative protocol to communicate. Many applications will try alternative techniques to avoid blocking.

As such great care and expertise is involved in writing signatures to achieve the desired effect.

Note: Application Control Lite, while powerful, can be difficult, time consuming, and dangerous to configure correctly. Application Control is recommended for most users as it comes preloaded with hundreds of maintained and current behavioral signatures and a commercial third-party application identification engine.

Warning: In previous and older versions of NG Firewall, there was a default signature set. However, enabling block on some of the default signatures caused false positives and blocked legitimate network traffic. Despite big warnings in the user interface and help documentation, we found users often misconfigured Application Control Lite anyway and experienced network problems as a result. To avoid this issue there are now no default signatures. If you have read this warning and understand that misconfiguring Application Control Lite **will cause major network connectivity issues,** then you can download the original list of signatures

[here](#) and import them into Signatures.

Reporting

The Reports tab provides a view of all reports and events for all traffic handled by Application Control Lite.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Application Control Lite Summary	A summary of Application Control Lite actions.
Detection Statistics	The number of logged and blocked sessions over time.
Top Blocked Protocols	The top blocked sessions by protocol.
Top Logged Protocols	The top logged sessions by protocol.
Top Blocked Hosts	The top blocked sessions by host.
Top Logged Hosts	The top logged sessions by host.
Top Blocked Users	The top blocked sessions by user.
Top Logged Users	The top logged sessions by user.
All Events	All sessions scanned by Application Control Lite.
Blocked Events	All sessions matching an application signature and blocked.

The tables queried to render these reports:

- [Sessions](#)

- Session Minutes

Related Topics

Application Control

8.3.1   Application Control Lite Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Application Control Lite.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Application Control Lite Summary	A summary of Application Control Lite actions.
Detection Statistics	The number of logged and blocked sessions over time.
Top Blocked Protocols	The top blocked sessions by protocol.
Top Logged Protocols	The top logged sessions by protocol.
Top Blocked Hosts	The top blocked sessions by host.
Top Logged Hosts	The top logged sessions by host.
Top Blocked Users	The top blocked sessions by user.
Top Logged Users	The top logged sessions by user.
All Events	All sessions scanned by Application Control Lite.
Blocked Events	All sessions matching an application signature and blocked.

The tables queried to render these reports:

- Sessions
- Session Minutes

8.4   SSL Inspector

The SSL Inspector is a special application that allows other NG Firewall applications that process HTTP traffic to also process encrypted HTTPS traffic and applications that process SMTP to also process SMTP over SSL. It does this by performing man-in-the-middle decryption and encryption of SSL traffic, passing the unencrypted traffic through the NG Firewall server for inspection by other applications and services.

About SSL Inspector

When a client makes an HTTPS request, the Inspector first initiates a secure SSL connection with the external server on behalf of the client. While this session is being established, the inspector captures information about the server SSL certificate. Once the server session is active, the Inspector uses the details from the server certificate to create a new certificate that will be used to encrypt the session between the inspector and the client. This certificate is generated or loaded on the fly, and is created using the same subject details contained in the actual server certificate. The certificate is then signed by the internal CA on the NG Firewall server, and is used to establish a secure connection between the inspector and the client. Creating the certificate this way is necessary to eliminate security warnings on the client, but it does require a few extra steps to properly configure the client computers and devices on your network. See the SSL Certificates section below for details.

SSL Certificates

SSL Certificates serve two primary purposes. They allow traffic between the client and server to be encrypted, and they allow the client to validate the authenticity of the server. There are two main ways the client checks the authenticity of the server certificate. The first is by validating the server certificate to ensure it has been issued or signed by a known and trusted third party certificate authority. Once that trust has been established, the client checks the server name portion of the target URL to ensure it matches the server name registered in the certificate presented by the server. If either of these checks fail, the client will typically display a warning, indicating that the security of the connection may be compromised.

When the NG Firewall server is initially installed, a default Certificate Authority is created automatically and used to sign the man-in-the-middle certificates created by the SSL Inspector. To view or make changes to the internal Certificate Authority, check out Certificates tab of the Config/Administration page.

Config > **Administration** > **Certificates**

Client Configuration

For the client authenticity checks to be successful, the client must be configured to trust the root certificate used by the NG Firewall server to sign the man-in-the-middle certificates described above. To configure clients, you must first use the Download Root Certificate button located on the Configuration tab of SSL Inspector Settings page to download the root certificate. You must then install this certificate in correct the location on the client.

Another way to download the root certificate is to simply access a special URL using the IP address of the NG Firewall server:

http://yourserver/cert

Simply replace *0.0.0.0* with the IP address of your NG Firewall server. This method is especially useful when using mobile devices. For example, accessing this URL on an iPad or iPhone will download and display the certificate, and provide an option to install and trust the certificate directly on the device.

Below are basic instructions for installing the root certificate on some common client platforms. If yours is not listed, or you have any difficulty, consult the reference material for the target platform for further information.

Internet Explorer or Google Chrome on Microsoft Windows

Follow the below steps:

1. Log into the NG Firewall server running SSL Inspector.
2. Go to **Config > Administration > Certificates** and download the certificate using the "**Download Root Certificate Authority (CA)**" button.
3. Copy the `root_authority.crt` you just downloaded to the Windows client computer.
4. From a command prompt, or from Start/Run, run the command "`certmgr.msc`".
5. Open the "Trusted Root Certification Authorities" tree in the panel on the left.
6. Right click on "**Certificates**" and select **All Tasks > Import** .
7. Proceed with the Certificate Import Wizard, selecting the the `root_authority.crt` file.

Firefox on Microsoft Windows

Follow the below steps:

1. Log into the NG Firewall server running SSL Inspector.
2. Go to **Config > Administration > Certificates** and download the certificate using the "**Download Root Certificate Authority (CA)**" button.
3. Copy the `root_authority.crt` you just downloaded to the Windows client computer.
4. Launch Firefox
5. From the Tools menu, go to **Options > Privacy & Security**.
6. Click the Import button and select the `root_authority.crt` file.
7. Enable the "Trust this CA to identify websites" checkbox and click the **OK** button.

Opera on Microsoft Windows

Follow the below steps:

1. Log into the NG Firewall server running SSL Inspector.
2. Go to **Config > Administration > Certificates** and download the certificate using the **Download Root Certificate Authority (CA)** button.
3. Copy the `root_authority.crt` you just downloaded to the Windows client computer.
4. Launch Opera
5. From the Tools menu, go to **Preferences > Advanced > Security** and click **Manage Certificates**.
6. Select the **Authorities** tab, click **Import**, and select the `root_authority.crt` file.
7. Click Install and click **OK** when asked if you are sure you want to trust the certificate.

Group Policy Distribution

If you have a fully deployed and implemented Active Directory infrastructure, you can leverage the Group Policy model to distribute the NG Firewall root certificate to all of your client computers. This is way outside our own area of expertise, so we can't provide much help or assistance, but we have compiled links to some TechNet articles with instructions for several common versions of Windows Server.

Windows Server 2003

Windows Server 2008

Windows Server 2012

Settings

This section describes the different settings and configuration options available for SSL Inspector.

Status

This displays the current status and some statistics.
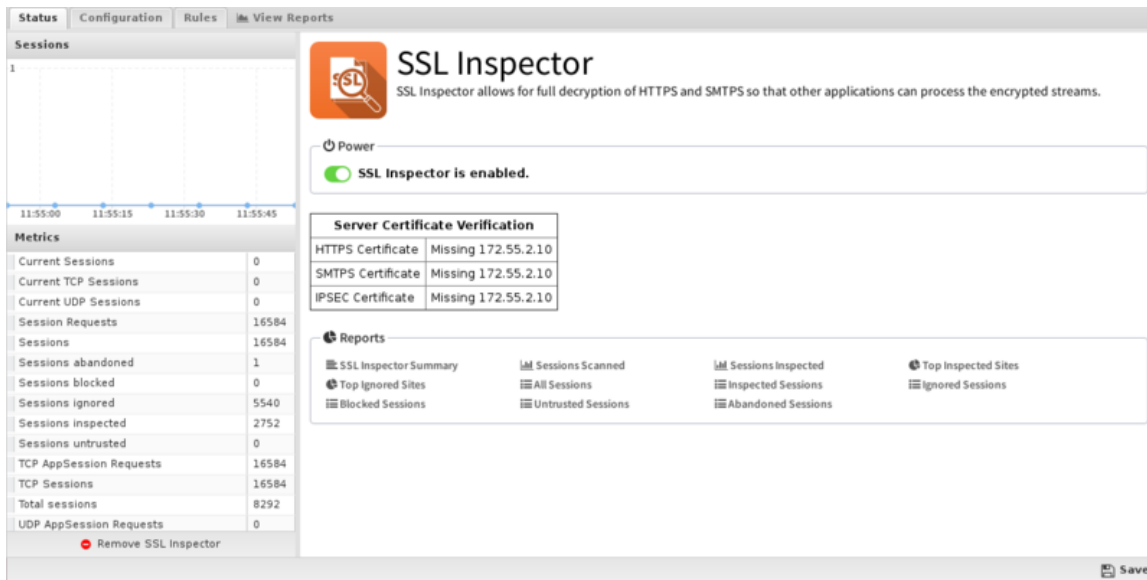
Figure 58: SSL Inspector Status

Configuration

Execute the following configurations:

Download Root Certificate

As described above, client computers and devices on your network need to be configured to trust the root certificate of the NG Firewall server. Clicking this button will allow you to download the root certificate. Once downloaded, you need to install it in the Trusted Authorities certificate store on your client computers and devices. Note that this is the same root certificate that can be downloaded from the **Config** > **Administration** > **Certificates** page. The download link is included on the SSL Inspector Configuration page for convenience.

Alternatively, you can download the certificate from a client system by navigating to http://yourserver/cert.

Enable SMTPS Traffic Processing

This option is enabled by default, and allows the SSL Inspector to work cooperatively with the other applications that act on SMTP mail traffic. When enabled, port *25* mail sessions that use STARTTLS will be decrypted inbound, allowing the clear traffic to pass through all other applications, and the re-encrypted again before passing outbound.

Enable HTTPS Traffic Processing

This option is enabled by default, and allows the SSL Inspector to work cooperatively with the other applications that act on HTTP web traffic. When enabled, port *443* web sessions that use SSL/TLS will be decrypted inbound, allowing the clear traffic to pass through all other applications, and then re-encrypted again before passing outbound.

Block Invalid HTTPS Traffic

When processing a new **HTTPS** session, the first thing the inspector does is analyze the initial client request to see if it contains a valid SSL negotiation message. If not, by default the session will be ignored and the traffic will flow directly between the client and server with no inspection performed. By enabling this checkbox, you can change the default behavior and effectively block any port *443* traffic that does not contain a valid HTTPS signature.

Client/Server Connection Protocols

This section includes checkboxes for turning on and off the SSL and TLS protocols that will be used when negotiating secure HTTPS and SMTPS inbound and outbound connections. The client protocols are used when the server is communicating with a client. The server protocols are used when the server is communicating with a server.

- **SSLv2Hello** - This is really a legacy handshake protocol that is used between a client and server when deciding which encryption protocol to use. This means it's possible to enable **SSLv2Hello** and still have a **TLSv1.x** connection negotiated. While there are no known security issues, **we still recommend leaving this disabled** unless you specifically need this legacy support.
- **SSLv3** - This is an older protocol that is now deprecated since the discovery of the POODLE vulnerability. For that reason, we recommend this be disabled for maximum security.
- **TLSv1** - This is an older protocol that has some known weaknesses. These can be mitigated if the other side of the connection forces certain secure ciphers to be used. However, since this can't be guaranteed, best practice is to disable this protocol unless it is required to support connections with legacy clients or servers.
- **TLSv1.1** - This is a modern protocol that is generally regarded as secure and is used as a fallback for *1.2* or in older browsers.
- **TLSv1.2** - This is the most common and recommended TLS version.
- **TLSv1.3**- This is the most recent version of the TLS protocol and offers the highest security but may have issues with some websites.

Trust All Server Certificates

Normally, when establishing an SSL connection with an external web server, the inspector will authenticate the server certificate against a standard list of trusted certificate authorities. If this trust cannot be established, the inspector will end the session. By enabling this checkbox, you can force the inspector to blindly trust all external server certificates.

Please note that we **DO NOT** recommend running with this option enabled, as it exposes all HTTPS traffic to significant security risks.

The standard list of trusted certificates used by NG Firewall is generated from the standard ca-certificates package. It includes, among others, certificate authorities used by Mozilla's browsers. Please note that Edge Threat Management staff can neither confirm nor deny whether the certificate authorities whose certificates are included in this list have in any way been audited for trustworthiness or *RFC 3647* compliance. Full responsibility to assess them belongs to the local system administrator.

Upload Trusted Certificate

Note: This setting applies to all policies when when using SSL Inspector with Policy Manager.

The inspector emulates a web browser when it makes outbound connections to external web servers. Just like a web browser, it must verify the authenticity of the server certificate before it will trust the connection and allow traffic to flow freely. As mentioned above, the inspector uses a standard list of known certificate authorities to validate server certificates. However, it's also possible you have servers in your network that use certificates that can't be authenticated this way. Perhaps you have your own certificate authority, or use self-signed certificates. Whatever the reason, you can use this section of the configuration page to upload additional certificates that you want the inspector to trust.
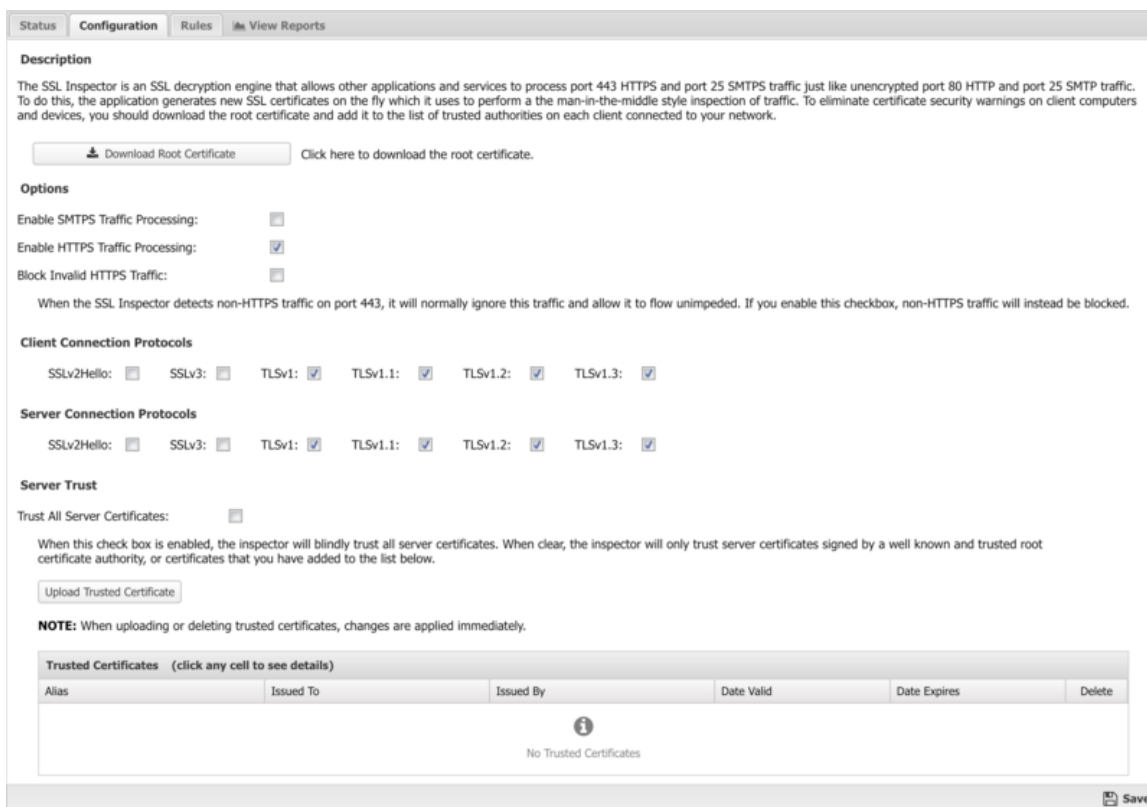
Figure 59: SSL Inspector Configuration

Rules

The **Rules** tab allows you to specify explicit rules to Inspect or Ignore HTTPS traffic that crosses the NG Firewall. By default, many common HTTPS sites (Google, YouTube, Yahoo, and so on) are inspected, but not all HTTPS. This provides a safe default which provides HTTPS inspection on those sites without interfering with other HTTPS communications. It can easily be configured to inspect all HTTPS by enabling the "Inspect All Traffic" rule.

The Rules describes how rules work and how they are configured. SSL Inspector uses rules to determine if it should inspect or ignore traffic for the specific session.

In addition to all the common rule types, there are three that are unique to the SSL Inspector, and these can be very useful for ignoring traffic that you don't want to inspect, or that isn't compatible with the SSL Inspector.

HTTPS: SNI Hostname

Most web browsers and many client applications include the destination hostname in the initial packet of an HTTPS session. The mechanism used is called the Server Name Indication, or the SNI extension to the TLS protocol. The main purpose is to allow a single web server to host multiple secure web sites. By analyzing the SNI hostname in the client request, the server can decide which SSL certificate to use for encrypting the session. This extension is necessary because the encryption must be established long before the server ever sees the HTTP request, and by then it would be too late to use a different certificate.

Creating ignore rules based on the SNI hostname is an effective way to have the SSL Inspector ignore incompatible traffic. A prime example is the default rule for Microsoft Update. The Microsoft Update client checks the server certificate to ensure it was signed by a specific authority. Since it doesn't trust the Root Authority the SSL Inspector uses to generate certificates on-the-fly, Microsoft Update will fail with an error. The default rule allows this traffic to be detected and ignored, allowing Microsoft Update to work properly.

HTTPS: Certificate Subject and **HTTPS: Certificate Issuer**

These two rule conditions are useful when dealing with client applications that don't use SNI, and aren't compatible with SSL Inspector. An excellent example is the Dropbox client utility for which there is also a default rule. Like Microsoft Update, the Dropbox client will reject SSL certificates that it doesn't explicitly trust.

Using either of these rule conditions, you can match traffic on any portion of the Subject or Issuer Distinguished Name (DN) included in the server certificate. In both cases, the information in the match string includes the standard information fields commonly stored within the SSL certificates, such as CN (common name), C (country), ST (state), L (locality), O (organization), and OU (organizational unit). Each of these are appended to the match string and separated by commas. Note that not all fields are required in all certificates, and some certificates may have others not listed. The order they occur in the match string is also not guaranteed.

The Subject DN generally includes information about the company to which the certificate was issued. Here is an example Certificate Subject:

```
CN=*.dropbox.com, O="Dropbox, Inc.", L=San Francisco, ST=California, C=US
```

The Issuer DN generally includes information about the company that issued and authenticated the certificate. Here is an example Certificate Issuer:

```
CN=Thawte SSL CA, O="Thawte, Inc.", C=US
```

Rule Actions

- **Inspect**: Causes the traffic which matched the rule to be decrypted and passed along to other applications and service for further inspection, classification, and possible action.
- **Ignore**: Causes the traffic which matched the rule to be ignored the SSL Inspector.

Figure 60: SSL Inspector Rules

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by HTTPS Inspector.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
SSL Inspector Summary	A summary of SSL Inspector actions.
Sessions Scanned	The amount of SSL sessions over time.
Sessions Inspected	The amount of inspected SSL sessions over time.
Top Inspected Sites	The number of inspected sessions grouped by site.
Top Ignored Sites	The number of ignored sessions grouped by site.
All Sessions	All sessions detected by SSL Inspector.
Inspected Sessions	Events where traffic was fully processed by the inspector, and all traffic was passed through all the other applications and services.
Ignored Sessions	Events where traffic was not or could not be inspected, so the traffic was completely ignored and not analyzed by any applications or services.
Blocked Sessions	Events where traffic was blocked because it did not contain a valid SSL request, and the Block Invalid Traffic option was enabled.
Untrusted Sessions	Events where traffic was blocked because the server certificate could not be authenticated.
Abandoned Sessions	Events where traffic was blocked due to an underlying problems with the SSL session.

The tables queried to render these reports:

- Database Schema

Status

The status of the session that generated the event.

- INSPECTED means the session was fully processed by the inspector, and all traffic was passed through all the other applications and services.
- IGNORED means the session was not or could not be inspected, so the traffic was completely ignored and not analyzed by any applications or services.
- BLOCKED means the traffic was blocked because it did not contain a valid **HTTPS** request, and the Block Invalid Traffic option was enabled.
- UNTRUSTED means the traffic was blocked because the server certificate could not be authenticated.
- ABANDONED means the connection failed because an an underlying SSL connection problem. Usually that the client abandoned the connection because the certificate was not trusted.

Detail

Extra details about the session, with the exact content dependent on the event status.

For INSPECTED and UNTRUSTED sessions, this field will include the SNI hostname extracted from the initial message sent from the client to the server. If the SNI information is not available, the server IP address will be used instead.

For BLOCKED or IGNORED sessions, this field will contain the description of the rule that matched and was applied to the session.

For ABANDONED sessions, detail will usually record information about the error that caused inspection to fail. For SSL exceptions, this will include Client or Server to indicate the session endpoint for which traffic was being processed. It will also include Encrypt or Decrypt to indicate the state of traffic inspection when the exception occurred. If available, the SSL error message will also be included. The following table lists the most common error messages and detailed information about each one.

Table 7: SSL Exception Messages

SSL Exception Message	Description
unexpected_message	An inappropriate message was received. This alert is always fatal and should never be observed in communication between proper implementations.
bad_record_mac	This alert is returned if a record is received with an incorrect MAC. This alert also MUST be returned if an alert is sent because a **TLSCiphertext** decrypted in an invalid way: either it wasn't an even multiple of the block length, or its padding values, when checked, weren't correct. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).
decryption_failed	This alert was used in some earlier versions of TLS, and may have permitted certain attacks against the CBC mode [CBCATT]. It MUST NOT be sent by compliant implementations.
record_overflow	A TLSCiphertext record was received that had a length more than 2^14+2048 bytes, or a record decrypted to a **TLSCompressed** record with more than 2^14+1024 bytes. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).
decompression_failure	The decompression function received improper input (e.g., data that would expand to excessive length). This message is always fatal and should never be observed in communication between proper implementations.
handshake_failure	Reception of a **handshake_failure** alert message indicates that the sender was unable to negotiate an acceptable set of security parameters given the options available. This is a fatal error.
no_certificate	This alert was used in **SSLv3** but not any version of TLS. It MUST NOT be sent by compliant implementations.
bad_certificate	A certificate was corrupt, contained signatures that did not verify correctly, etc.
unsupported_certificate	A certificate was of an unsupported type.
certificate_revoked	A certificate was revoked by its signer.

SSL Exception Message	Description
certificate_expired	A certificate has expired or is not currently valid.
certificate_unknown	Some other (unspecified) issue arose in processing the certificate, rendering it unacceptable.
illegal_parameter	A field in the handshake was out of range or inconsistent with other fields. This message is always fatal.
unknown_ca	A valid certificate chain or partial chain was received, but the certificate was not accepted because the CA certificate could not be located or couldn't be matched with a known, trusted CA. This message is always fatal.
access_denied	A valid certificate was received, but when access control was applied, the sender decided not to proceed with negotiation. This message is always fatal.
decode_error	A message could not be decoded because some field was out of the specified range or the length of the message was incorrect. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).
decrypt_error	A handshake cryptographic operation failed, including being unable to correctly verify a signature or validate a Finished message. This message is always fatal.
export_restriction	This alert was used in some earlier versions of TLS. It MUST NOT be sent by compliant implementations.
protocol_version	The protocol version the client has attempted to negotiate is recognized but not supported. (For example, old protocol versions might be avoided for security reasons.) This message is always fatal.
insufficient_security	Returned instead of **handshake_failure** when a negotiation has failed specifically because the server requires ciphers more secure than those supported by the client. This message is always fatal.
internal_error	An internal error unrelated to the peer or the correctness of the protocol (such as a memory allocation failure) makes it impossible to continue. This message is always fatal.
user_canceled	This handshake is being canceled for some reason unrelated to a protocol failure. If the user cancels an operation after the handshake is complete, just closing the connection by sending a **close_notify** is more appropriate. This alert should be followed by a **close_notify**. This message is generally a warning.
no_renegotiation	Sent by the client in response to a hello request or by the server in response to a client hello after initial handshaking. Either of these would normally lead to renegotiation; when that is not appropriate, the recipient should respond with this alert. At that point, the original requester can decide whether to proceed with the connection. One case where this would be appropriate is where a server has spawned a process to satisfy a request; the process might receive security parameters (key length, authentication, etc.) at startup, and it might be difficult to communicate changes to these parameters after that point. This message is always a warning.

SSL Exception Message	Description
unsupported_extension	sent by clients that receive an extended server hello containing an extension that they did not put in the corresponding client hello. This message is always fatal.

8.4.1 SSL Inspector Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by HTTPS Inspector.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
SSL Inspector Summary	A summary of SSL Inspector actions.
Sessions Scanned	The amount of SSL sessions over time.
Sessions Inspected	The amount of inspected SSL sessions over time.
Top Inspected Sites	The number of inspected sessions grouped by site.
Top Ignored Sites	The number of ignored sessions grouped by site.
All Sessions	All sessions detected by SSL Inspector.
Inspected Sessions	Events where traffic was fully processed by the inspector, and all traffic was passed through all the other applications and services.
Ignored Sessions	Events where traffic was not or could not be inspected, so the traffic was completely ignored and not analyzed by any applications or services.
Blocked Sessions	Events where traffic was blocked because it did not contain a valid SSL request, and the Block Invalid Traffic option was enabled.
Untrusted Sessions	Events where traffic was blocked because the server certificate could not be authenticated.
Abandoned Sessions	Events where traffic was blocked due to an underlying problems with the SSL session.

The tables queried to render these reports:

- Database Schema

Status

The status of the session that generated the event.

- INSPECTED means the session was fully processed by the inspector, and all traffic was passed through all the other applications and services.
- IGNORED means the session was not or could not be inspected, so the traffic was completely ignored and not analyzed by any applications or services.
- BLOCKED means the traffic was blocked because it did not contain a valid HTTPS request, and the Block Invalid Traffic option was enabled.
- UNTRUSTED means the traffic was blocked because the server certificate could not be authenticated.
- ABANDONED means the connection failed because an an underlying SSL connection problem. Usually that the client abandoned the connection because the certificate was not trusted.

Detail

Extra details about the session, with the exact content dependent on the event status.

For INSPECTED and UNTRUSTED sessions, this field will include the SNI hostname extracted from the initial message sent from the client to the server. If the SNI information is not available, the server IP address will be used instead.

For BLOCKED or IGNORED sessions, this field will contain the description of the rule that matched and was applied to the session.

For ABANDONED sessions, detail will usually record information about the error that caused inspection to fail. For SSL exceptions, this will include Client or Server to indicate the session endpoint for which traffic was being processed. It will also include Encrypt or Decrypt to indicate the state of traffic inspection when the exception occurred. If available, the SSL error message will also be included. The following table lists the most common error messages and detailed information about each one.

Table 8: SSL Exception Messages

SSL Exception Messages	Description
unexpected_message	An inappropriate message was received. This alert is always fatal and should never be observed in communication between proper implementations.
bad_record_mac	This alert is returned if a record is received with an incorrect MAC. This alert also MUST be returned if an alert is sent because a **TLSCiphertext** decrypted in an invalid way: either it wasn't an even multiple of the block length, or its padding values, when checked, weren't correct. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).
decryption_failed	This alert was used in some earlier versions of TLS, and may have permitted certain attacks against the CBC mode [CBCATT]. It MUST NOT be sent by compliant implementations.
record_overflow	A **TLSCiphertext** record was received that had a length more than 2^14+2048 bytes, or a record decrypted to a **TLSCompressed** record with more than 2^14+1024 bytes. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).

SSL Exception Messages	Description
decompression_failure	The decompression function received improper input (e.g., data that would expand to excessive length). This message is always fatal and should never be observed in communication between proper implementations.
handshake_failure	Reception of a **handshake_failure** alert message indicates that the sender was unable to negotiate an acceptable set of security parameters given the options available. This is a fatal error.
no_certificate	This alert was used in SSLv3 but not any version of TLS. It MUST NOT be sent by compliant implementations.
bad_certificate	A certificate was corrupt, contained signatures that did not verify correctly, etc.
unsupported_certificate	A certificate was of an unsupported type.
certificate_revoked	A certificate was revoked by its signer.
certificate_expired	A certificate has expired or is not currently valid.
certificate_unknown	Some other (unspecified) issue arose in processing the certificate, rendering it unacceptable.
illegal_parameter	A field in the handshake was out of range or inconsistent with other fields. This message is always fatal.
unknown_ca	A valid certificate chain or partial chain was received, but the certificate was not accepted because the CA certificate could not be located or couldn't be matched with a known, trusted CA. This message is always fatal.
access_denied	A valid certificate was received, but when access control was applied, the sender decided not to proceed with negotiation. This message is always fatal.
decode_error	A message could not be decoded because some field was out of the specified range or the length of the message was incorrect. This message is always fatal and should never be observed in communication between proper implementations (except when messages were corrupted in the network).
decrypt_error	A handshake cryptographic operation failed, including being unable to correctly verify a signature or validate a Finished message. This message is always fatal.
export_restriction	This alert was used in some earlier versions of TLS. It MUST NOT be sent by compliant implementations.
protocol_version	The protocol version the client has attempted to negotiate is recognized but not supported. (For example, old protocol versions might be avoided for security reasons.) This message is always fatal.
insufficient_security	Returned instead of **handshake_failure** when a negotiation has failed specifically because the server requires ciphers more secure than those supported by the client. This message is always fatal.
internal_error	An internal error unrelated to the peer or the correctness of the protocol (such as a memory allocation failure) makes it impossible to continue. This message is always fatal.

SSL Exception Messages	Description
user_canceled	This handshake is being canceled for some reason unrelated to a protocol failure. If the user cancels an operation after the handshake is complete, just closing the connection by sending a **close_notify** is more appropriate. This alert should be followed by a **close_notify**. This message is generally a warning.
no_renegotiation	Sent by the client in response to a hello request or by the server in response to a client hello after initial handshaking. Either of these would normally lead to renegotiation; when that is not appropriate, the recipient should respond with this alert. At that point, the original requester can decide whether to proceed with the connection. One case where this would be appropriate is where a server has spawned a process to satisfy a request; the process might receive security parameters (key length, authentication, etc.) at startup, and it might be difficult to communicate changes to these parameters after that point. This message is always a warning.
unsupported_extension	sent by clients that receive an extended server hello containing an extension that they did not put in the corresponding client hello. This message is always fatal.

Related Topics

Report Viewer

Reports

8.5 Spam Blocker

Spam Blocker is an intelligent email filter that identifies and handles spam (unsolicited bulk email). It leverages technology from the SpamAssassin project and improves upon it with the integration of a commercial spam engine. It can scan any email that is transported via SMTP.

Spam Blocker transparently scans email transported over SMTP on the way to your mail server (or outbound if configured so). It does not require any reconfiguration of your DNS MX records or the email server. Any SMTP traffic going through the NG Firewall server will be scanned.

Settings

This section reviews the different settings and configuration options available for Spam Blocker.

Status

This displays the current status and some statistics.

Email

These settings apply only to the SMTP protocol.

- **Scan SMTP**: This enables or disables SMTP scanning.

- **Strength**: If the Spam Score of a message is equal to or greater than this setting your chosen action will be taken on the message. Higher values make Spam Blocker **more** sensitive to spam.

- **Action**: The action taken on the message if the Spam Score is high enough.

> If set to **Mark**, "[Spam]..." will be prepended to the email subject line and it will be delivered. If set to **Pass**, the message will be delivered as originally sent. **Drop** will inform the sending server the mail was successfully delivered, but NG Firewall will drop the mail so it is never delivered. **Quarantine** will send the mail to users' email quarantine for them to release or delete as they see fit. For more information, refer to Quarantine.

- **Drop Super Spam**: If this option is enabled, any emails with a score greater than the Super spam score will be dropped.

- **Super Spam Score**: The score emails must reach to be dropped as Super Spam.

- **Advanced SMTP Configuration**

> ⚠ **Important:** The default values are the suggested values. Changing and customizing settings can easily and will likely cause Spam Blocker to perform less than optimally.

> - **Enable tarpitting**: This option enables Tarpit. If enabled, when an SMTP session is first caught Spam Blocker will check if the client IP is on a Domain Name System BLocklist (DNSBL). If it is, the session is rejected before the remote server is even allowed to send the email. This increases the capacity of a given server by quite a bit and can also save bandwidth, but it can increase false positives if the remote email server has mistakenly been put on a blacklist. **This setting will not increase spam scanning accuracy - it will actually decrease it as it will prevent valuable super-spam training data from reaching the spam engine.** By enabling this feature you are getting lower spam accuracy, but increased email scanning capacity.

> - **Enable greylisting**: This option enables greylisting. If enabled, each time a new sender tries to send mail to a specific receiver it will receive a "421 Please try again" error. The

second time will mail will be allowed. Greylisting will reduce spam because spammers often won't retry transmissions as they should, or the extra time delay will increase the chances of the spam engine properly identifying new spam waves. However, greylisting adds a delay to all legitimate email. This settings is not suggested for most sites because of the complications.

- **Add email headers**: When enabled, NG Firewall adds information about the Spam Score and the test run to get that score to the headers of the message.

- **Close connection on scan failure**: This option will close the connection if the scan fails so the message will be resent and retested. If disabled, a scan failure will allow the email to be delivered without being scanned.

- **Scan outbound (WAN) SMTP**: If unchecked, outbound mail (mail in sessions going out a WAN interface) will not be scanned. If checked, outbound mail will be scanned just like incoming mail.

- **Allow and ignore TLS sessions**: This option controls the allowance of TLS sessions. If unchecked (the default) the TLS advertisement (if present) is removed from the server advertisements and TLS is not allowed on any scanned sessions. If checked, the TLS advertisement is allowed and if the client initialized TLS the message will pass through completely unscanned, even if it is spam.

- **CPU Load Limit**: If your CPU Load exceeds this number incoming connections are stopped until the load decreases. This is specified so that spam scanning can not monopolize the server resources.

- **Concurrent Scan Limit**: This is the maximum number of messages that can be scanned at the same time. This is specified so that spam scanning can not monopolize the server resources.

- **Message Size Limit**: This option allows you to change the maximum size of a message that will be scanned for spam. The default maximum size is 256KB. Spam will typically be much smaller, as spammers rely on the sheer number of messages sent.

  Note: This does not control the message size limit of messages passed through NG Firewall. This does not effect the maximum size of message your server will accept, only the limit on the size of message that will be checked for spam.

Status	Email	▥ View Reports

☑ **Scan SMTP**

Strength:	Medium (Threshold: 4.3)	⌄
Action:	Quarantine	⌄

☑ **Drop Super Spam**

Super Spam Threshold: 20 ⌄

⌄ **Advanced SMTP Configuration**

🖫 Save

Reporting

The Reports tab provides a view of all reports and events for all traffic handled by Spam Blocker.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Spam Blocker Summary	A summary of spam blocking actions for email activity.
Email Usage (all)	The amount of scanned, clean, and spam email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (clean)	The amount of clean email over time.
Email Usage (spam)	The amount of spam email over time.
Spam Ratio	The ratio of spam (true) to ham (false)
Top Spam Recipients	The number of email addresses with spam.
Top Spam Sender Addresses	The number of IP addresses sending spam.
All Email Events	All emails scanned by Spam Blocker.
All Spam Events	All emails marked as Spam.
Quarantined Events	All emails marked as Spam and quarantined.
Tarpit Events	All email sessions that were tarpitted.

The tables queried to render these reports:

- Sessions
- Session Minutes

Related Topics

- Spam Blocker Lite

8.5.1 Spam Blocker Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Spam Blocker.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Spam Blocker Summary	A summary of spam blocking actions for email activity.
Email Usage (all)	The amount of scanned, clean, and spam email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (clean)	The amount of clean email over time.
Email Usage (spam)	The amount of spam email over time.
Spam Ratio	The ratio of spam (true) to ham (false)
Top Spam Recipients	The number of email addresses with spam.
Top Spam Sender Addresses	The number of IP addresses sending spam.
All Email Events	All emails scanned by Spam Blocker.
All Spam Events	All emails marked as Spam.
Quarantined Events	All emails marked as Spam and quarantined.
Tarpit Events	All email sessions that were tarpitted.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Spam Blocker

8.6   Spam Blocker Lite

Spam Blocker Lite is an intelligent email filter that identifies spam (unsolicited bulk email). It leverages technology from the SpamAssassin project. It can scan any email that is transported via SMTP.

Spam Blocker Lite transparently scans email transported over SMTP on the way to your mail server (or outbound if configured so). It does not require any reconfiguration of DNS MX records or any email servers. Any SMTP traffic going through the NG Firewall server will be scanned.

Settings

This section reviews the different settings and configuration options available for Spam Blocker Lite.

Status

This displays the current status and some statistics.

Email

These settings apply only to the SMTP protocol.

- **Scan SMTP**: This enables or disables SMTP scanning.
- **Strength**: If the Spam Score of a message is equal to or greater than this setting your chosen action will be taken on the message. Higher values make Spam Blocker *more* sensitive to spam.
- **Action**: The action taken on the message if the Spam Score is high enough.

  If set to **Mark**, "[Spam]..." will be prepended to the email subject line and it will be delivered. If set to **Pass**, the message will be delivered as originally sent. **Drop** will inform the sending server the mail was successfully delivered, but NG Firewall will drop the mail so it is never delivered. **Quarantine** will send the mail to users' email quarantine for them to release or delete as they see fit. For more information, refer to Quarantine.

- **Drop Super Spam**: If this option is enabled, any emails with a score greater than the Super spam score will be dropped.
- **Super Spam Score**: The score emails must reach to be dropped as Super Spam.
- **Advanced SMTP Configuration**:

  IMPORTANT: The default values are the suggested values. Changing and customizing settings can easily and will likely cause Spam Blocker to perform less than optimally.

  - **Enable tarpitting**: This option enables Tarpit. If enabled, when an SMTP session is first caught Spam Blocker will check if the client IP is on a Domain Name System BLocklist (DNSBL). If it is, the session is rejected before the remote server is even allowed to send the email. This increases the capacity of a given server by quite a bit and can also save bandwidth, but it can increase false positives if the remote email server has mistakenly been put on a blacklist. **This setting will not increase spam scanning accuracy - it will actually decrease it as it will prevent valuable super-spam training data from reaching the spam engine.** By enabling this feature you are getting lower spam accuracy, but increased email scanning capacity.

  - **Add email headers**: When enabled, NG Firewall adds information about the Spam Score and the test run to get that score to the headers of the message.

- **Close connection on scan failure**: This option will close the connection if the scan fails so the message will be resent and retested. If disabled, a scan failure will allow the email to be delivered without being scanned.

- **Scan outbound (WAN) SMTP**: If unchecked, outbound mail (mail in sessions going out a WAN interface) will not be scanned. If checked, outbound mail will be scanned just like incoming mail.

- **Allow and ignore TLS sessions**: This option controls the allowance of TLS sessions. If unchecked (the default) the TLS advertisement (if present) is removed from the server advertisements and TLS is not allowed on any scanned sessions. If checked, the TLS advertisement is allowed and if the client initialized TLS the message will pass through completely unscanned, even if it is spam.

- **CPU Load Limit**: If your CPU Load exceeds this number incoming connections are stopped until the load decreases. This is specified so that spam scanning can not monopolize the server resources.

- **Concurrent Scan Limit**: This is the maximum number of messages that can be scanned at the same time. This is specified so that spam scanning can not monopolize the server resources.

- **Message Size Limit**: This option allows you to change the maximum size of a message that will be scanned for spam. The default maximum size is 256KB. Spam will typically be much smaller, as spammers rely on the sheer number of messages sent.

  > **Note:** This does not control the message size limit of messages passed through NG Firewall. This does not effect the maximum size of message your server will accept, only the limit on the size of message that will be checked for spam.

Reporting

The Reports tab provides a view of all reports and events for all traffic handled by Spam Blocker Lite.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Spam Blocker Lite Summary	A summary of spam blocking actions for email activity.
Email Usage (all)	The amount of scanned, clean, and spam email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (clean)	The amount of clean email over time.
Email Usage (spam)	The amount of spam email over time.
Spam Ratio	The ratio of spam (true) to ham (false)
Top Spam Recipients	The number of email addresses with spam.
Top Spam Sender Addresses	The number of IP addresses sending spam.
All Email Events	All emails scanned by Spam Blocker.
All Spam Events	All emails marked as Spam.
Quarantined Events	All emails marked as Spam and quarantined.
Tarpit Events	All email sessions that were tarpitted.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Spam Blocker Lite Reports

Related Topics

- Spam Blocker

8.6.1   Spam Blocker Lite Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Spam Blocker Lite.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Spam Blocker Lite Summary	A summary of spam blocking actions for email activity.
Email Usage (all)	The amount of scanned, clean, and spam email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (clean)	The amount of clean email over time.

Report Entry	Description
Email Usage (spam)	The amount of spam email over time.
Spam Ratio	The ratio of spam (true) to ham (false)
Top Spam Recipients	The number of email addresses with spam.
Top Spam Sender Addresses	The number of IP addresses sending spam.
All Email Events	All emails scanned by Spam Blocker.
All Spam Events	All emails marked as Spam.
Quarantined Events	All emails marked as Spam and quarantined.
Tarpit Events	All email sessions that were tarpitted.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Spam Blocker Lite

8.7   Web Filter

Web Filter monitors HTTP and HTTPS traffic on your network to filter and log web activities and block inappropriate content. Web Filter also appeals to customers who require an added level of protection or are subject to regulations, for example Web Filter helps libraries comply with the Children's Internet Protection Act). Need to block Pornography or Hate Speech on your network? Web Filter is your answer.

About Web Filter

- **Real-time classification and updates**: When your users visit a site, NG Firewall sends the URL to the Webroot BrightCloud® to be categorized. When the data is returned, NG Firewall keeps a temporary local cache of the site and category to speed up the process the next time the URL is requested. This data is then used to flag or allow users access to the site they have requested, all without any appreciable increase in load time. If a site is not categorized upon request, it is autocategorized by our partners at Webroot and put into a queue to be verified by a human. Because this is done dynamically, new sites and updated URLs are allowed or flagged according to your settings without additional intervention, plus you have the option of requesting recategorization of sites.
- **HTTPS Filtering**: Web Filter has multiple techniques to deal with HTTPS, SSL-encrypted HTTP. HTTPS traffic is encrypted so only some information is visible and this information is used to categorize the session. More information on how this is down below.
- **Detailed categorization**: Web Filter offers 79 categories and tens of billions of URLs. The Web Filter database is over 100 times larger and more accurate. The abundance of categories means that you can narrow your scope - maybe you want to flag websites related to nudity, but allow sites dealing with Sexual Education.
- **Advanced features**: Force safe-search on search engines, filter and log user searches, restrict google domains, and more!

Traffic Flow

When scanning traffic, Web Filter evaluates the pass lists, block lists, categories, and rules at two distinct points of the HTTP transaction. The first evaluation happens after the request is received from the client and before it is forwarded to the server. The second is after the response is received from the server and before it is passed back to the client. This allows a high degree of filtering and control over both resources that are requested, and content that is returned in response.

HTTP Request

When evaluating HTTP requests, Web Filter applies the configured rules and lists in the following order:

1. A lookup is performed to determine the category for the requested site. The category is attached to the session for use by Web Filter as well as other applications.
2. The source IP of the request is checked against the Pass Clients list. If a match is found, the traffic is allowed.
3. The destination site of the request is checked against the Pass Sites list. If a match is found, the traffic is allowed.
4. If 'Restrict Google applications' is enabled, the appropriate header is added to the request using the 'Allowed Domains' that have been configured.
5. If 'Pass if referrer matches any Pass Sites' is enabled, the referrer is checked against the Pass Sites list. If a match is found, the traffic is allowed.
6. If the Unblock option is enabled, the destination site and the source IP of the request are checked against the unblock list. If a match is found, the traffic is allowed.
7. If 'Block pages from IP only hosts' is enabled, the request will be evaluated and blocked if the destination is an IP address.
8. The destination site of the request is checked against the Block Sites list. If a match is found, the traffic is blocked.
9. The traffic details are passed to the Rules list. If a match is found, the traffic is allowed, flagged, or blocked based on the options configured in the rule that was matched.
10. The category determined in **Step 1** is compared to the Categories list, and the traffic is allowed, flagged, or blocked based on the corresponding match. If the category could not be determined, the traffic is allowed.

HTTP Response

When evaluating HTTP responses, Web Filter applies the configured rules and lists in the following order:

1. The source IP of the request is checked against the Pass Clients list. If a match is found, the traffic is allowed.
2. The site from which the response was received is checked against the Pass Sites list. If a match is found, the traffic is allowed.
3. If the Unblock option is enabled, the site from which the response was received and the client IP are checked against the unblock list. If a match is found, the traffic is allowed.
4. The traffic details are passed to the Rules list. If a match is found, the traffic is allowed, flagged, or blocked based on the options configured in the rule that was matched.

Settings

This section reviews the different settings and configuration options available for Web Filter.

Status

This displays the current status and some statistics.
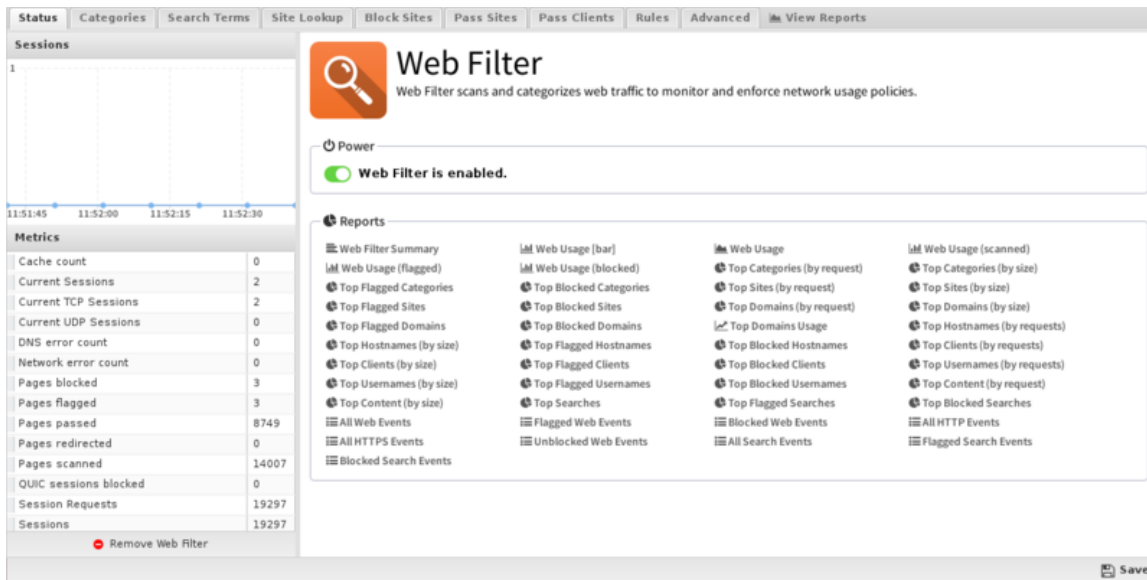
Figure 61: Web Filter Status

Categories

Categories enables you to customize which categories of sites will be blocked or flagged. Categories that are blocked will display a block page to the user; categories that are flagged will allow the user to access the site, but will be silently flagged as a violation for event logs and Reports. These block/flag actions operate the same way for all of the different Web Filter options.
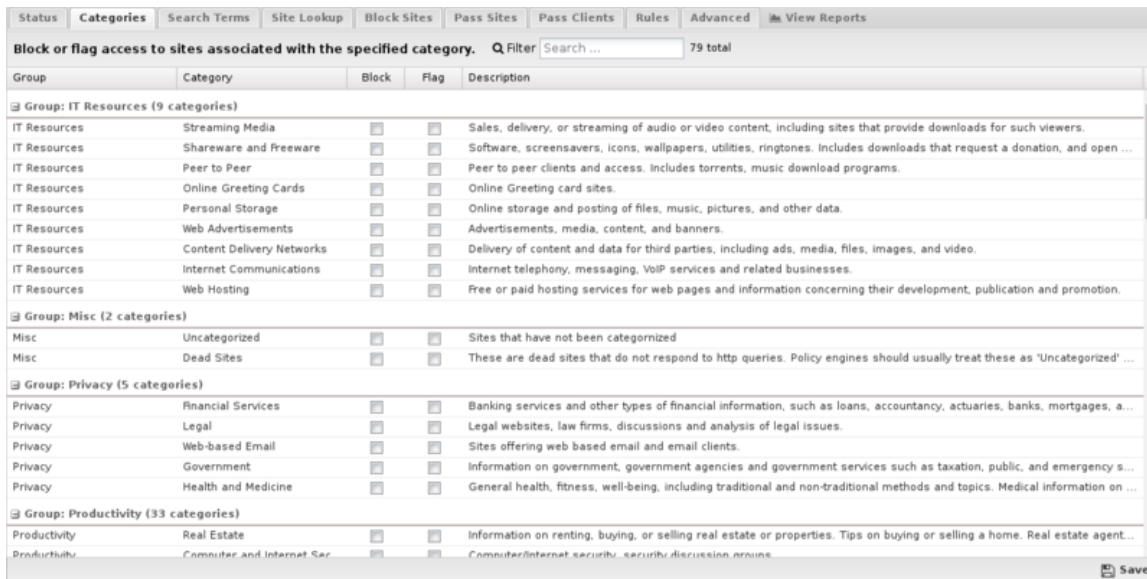
Figure 62: Web Filter Categories

Search Terms

Search Terms filtering enables you to flag or block specific search terms your users perform on popular search sites including Google, Bing, Ask, Yahoo, and YouTube. For example, if someone searches Google and includes the word "suicide", or they search for "twerking" videos on YouTube you can have these activities flagged or blocked.

Under Search Filter you can add terms you want to be blocked or flagged. Search Filter terms use the Glob Matcher syntax.
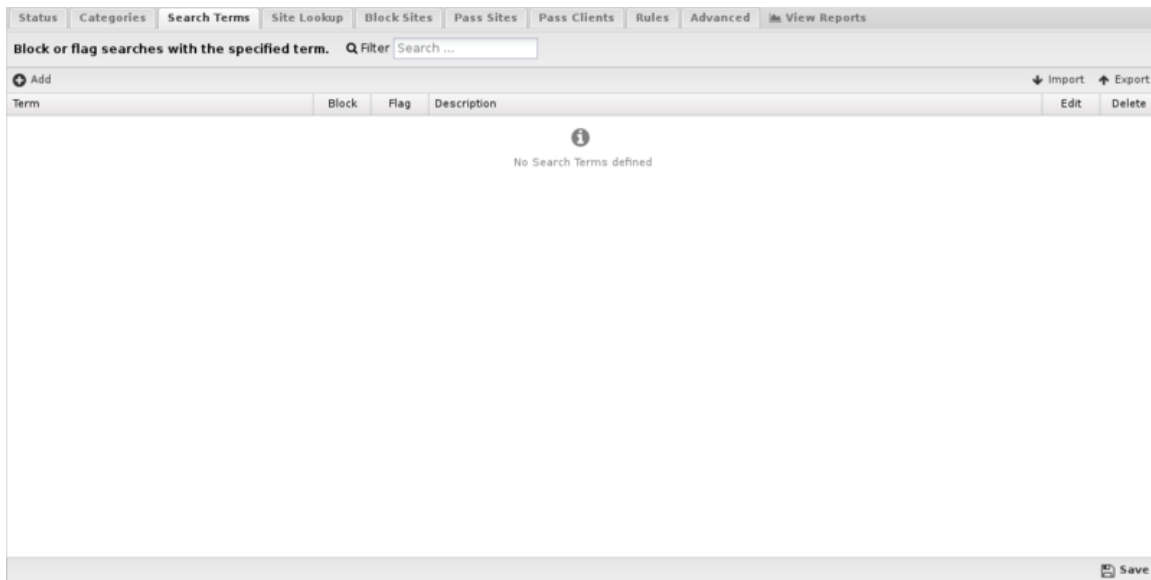
Figure 63: Web Filter Search Terms

In many cases you may have an existing list of search terms that you want to import. For example, you can find several banned words from Facebook, YouTube, Wordpress and other sources Full List of Bad Words and Top Swear Words Banned by Google Block Facebook, YouTube. The import feature enables you to import these lists from a comma separated or newline delimited file. You can also import in JSON format in case you are transferring a list from another NG Firewall.

Note: SSL Inspector must be installed and enabled to use search filtering.

Site Lookup

Site Lookup allows you to find the categorization of a URL. Clicking it brings up a dialog. In **Site URL** specify the URL to find and click **Search** to find the URL's categorization.

If you feel the current categorization is incorrect, check **Suggest a different category**, select a new category from the list, and click **Suggest** to submit the category change for consideration.

Note: This is only a suggestion - it will trigger the URL categorization provider to recategorize all categories for the URL.
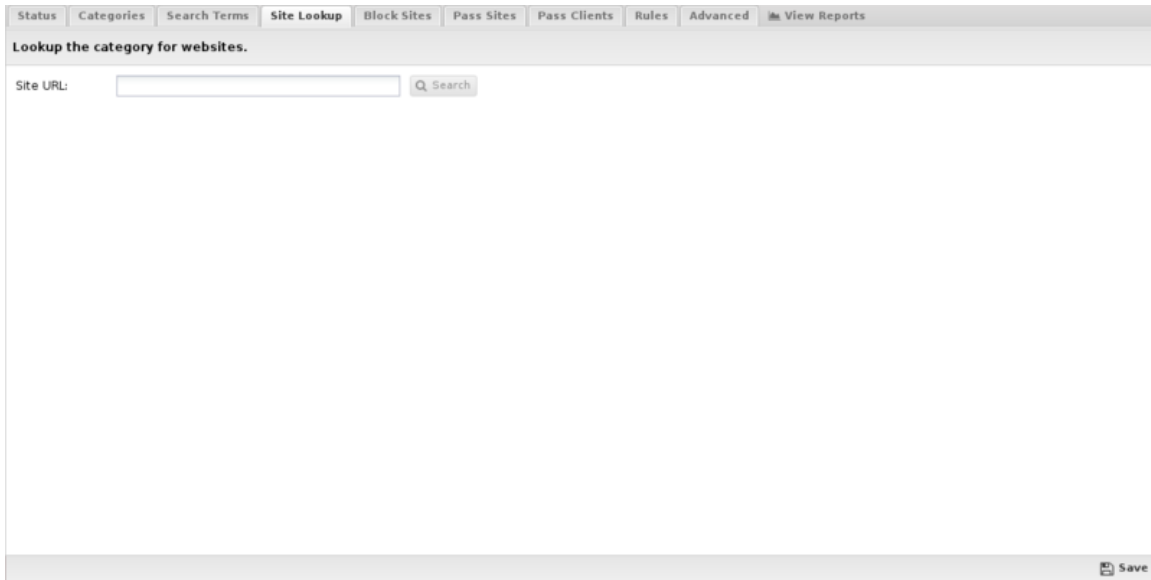
Figure 64: Web Filter Site Lookup

Block Sites

Under Block Sites you can add individual domain names you want to be blocked or flagged - just enter the domain name (e.g. youtube.com) and specify your chosen action. This list uses URL Matcher syntax.
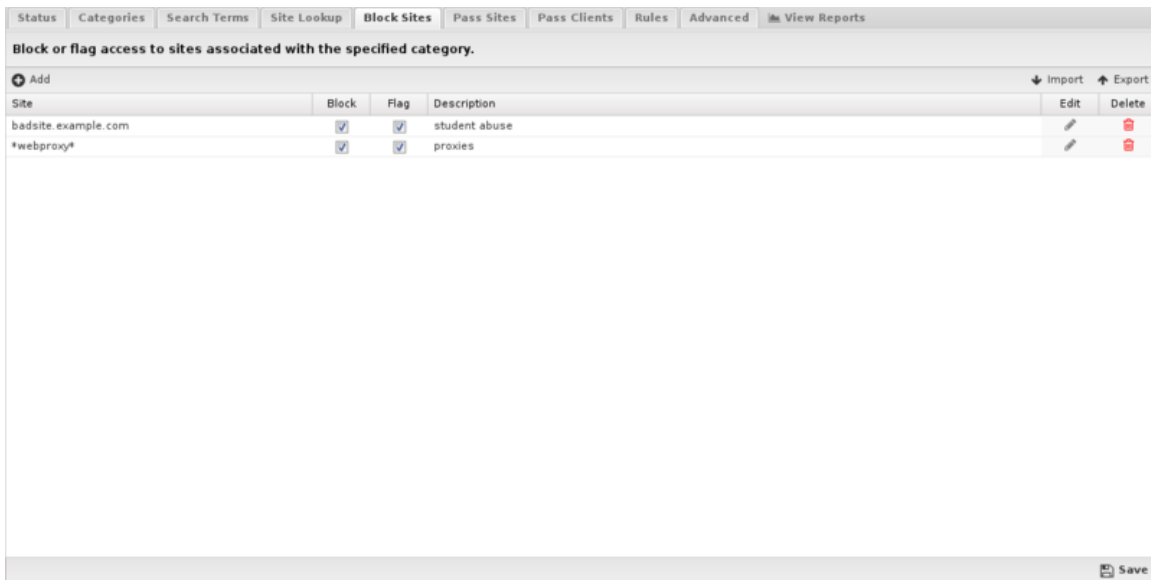
Figure 65: Web Filter Block Sites

Pass Sites

Pass Sites is used to pass content that would have otherwise been blocked. This can be useful for "unblocking" sites that you don't want blocked according to block settings. Any domains you add to the Passed Sites list will be allowed, even if blocked by category or by individual URL - just add the domain and save. Unchecking the pass option will allow the site to be blocked as if the entry was not present. This list uses URL Matcher syntax.

Figure 66: Web Filter Pass Sites

Pass Clients

If you add an IP address to this list, Web Filter will not block any traffic from that IP regardless of the blocked categories or sites. Just add the IP and save. Unchecking the pass option will have the block/ pass lists affect the user as if they were not entered into the Passed Client IPs list. This list uses IP Matcher syntax.

If you have a few users that need to completely bypass Web Filter controls, consider using pass lists. If you have users that simply need different Web Filter settings, you should set up a separate policy using Policy Manager. When using this feature, please remember that DHCP IPs can change, so you'll probably want to set up either a Static IP or a Static DHCP Lease for the machine in question.
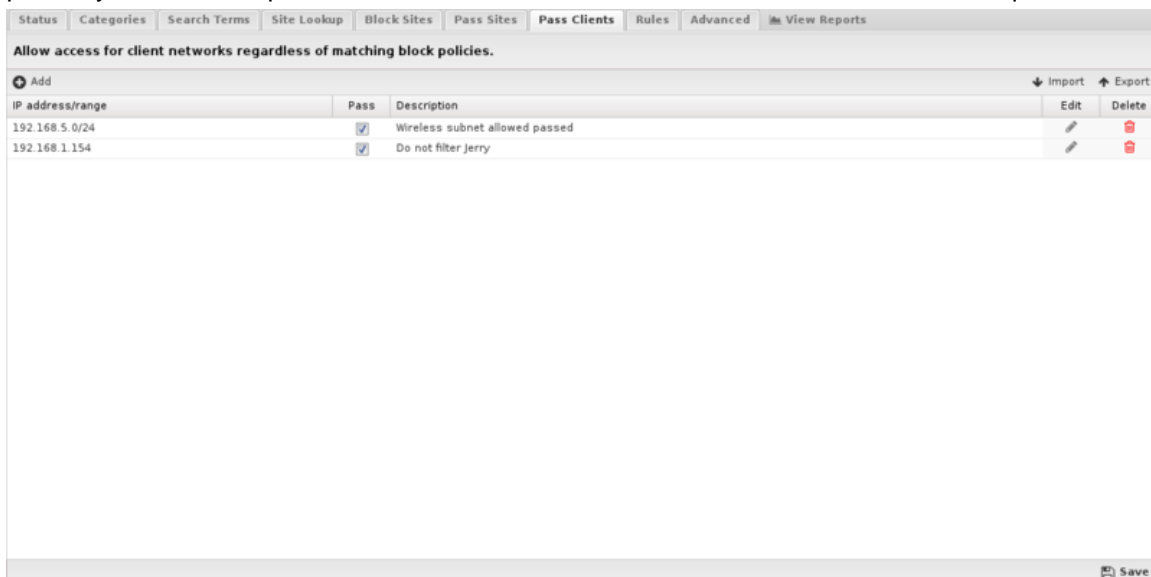
Figure 67: Web Filter Pass Clients

Rules

The **Rules** tab allows you to specify rules to Block or Flag traffic that passes through Web Filter.

The Rules describes how rules work and how they are configured. Web Filter uses rules to determine to block or flag the specific session. Flagging a session marks it in the logs for reviewing in the event logs or reports, but has no direct effect on the network traffic.
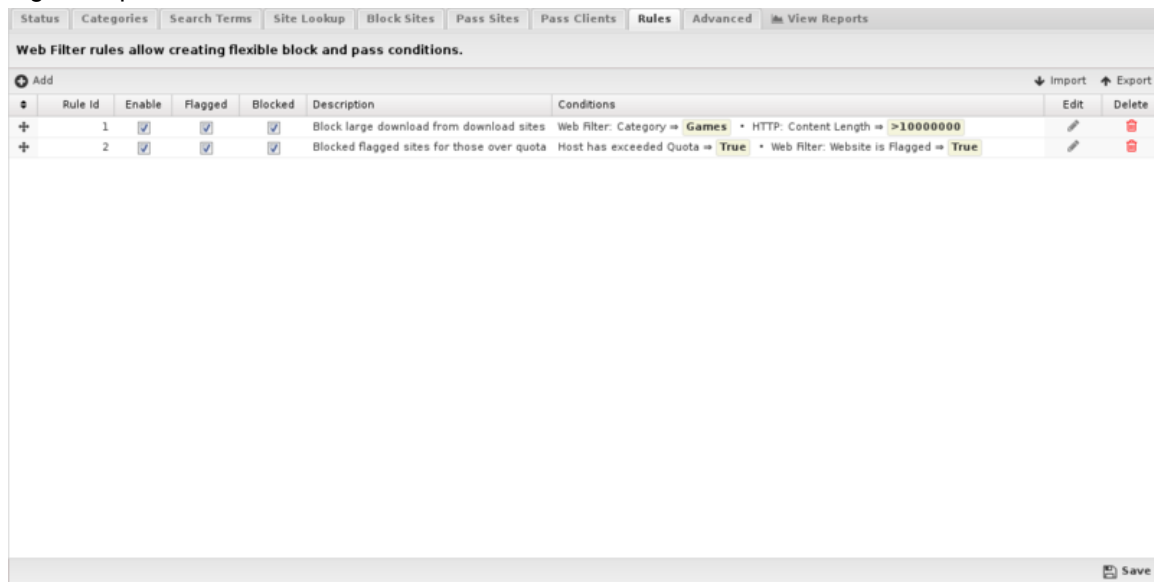
Figure 68: Web Filter Rules

Rule Actions

- **Flag**: Allows the traffic which matched the rule to flow, and flags the traffic for easier viewing in the event log.
- **Block**: Blocks the traffic which matched the rule.

Rule Types

In previous versions of Web Filter, there were dedicated lists for blocking certain file extensions or MIME types. This capability is still available using the more flexible filter rules. For blocking specific file extensions, you can create a rule with the condition **HTTP: Response File Extension** that has a comma separated list of the extensions to block in the Value field. For blocking MIME types, you would create a rule with the condition **HTTP: Content Type** that has a comma separated list of the content types to block in the Value field.

Below are tables that list the default file extensions and MIME types that were available in previous versions. Note that these lists are not exhaustive, but are included here as a reference, and to simplify creation of such rules via copy/paste of the values in the tables.

Extension	Category	Description
exe	executable	an executable file format
ocx	executable	an executable file format
dll	executable	an executable file format
cab	executable	an ActiveX executable file format
bin	executable	an executable file format
com	executable	an executable file format
jpg	image	an image file format
png	image	an image file format

Extension	Category	Description
gif	image	an image file format
jar	java	a Java file format
class	java	a Java file format
swf	flash	the flash file format
mp3	audio	an audio file format
wav	audio	an audio file format
wmf	audio	an audio file format
mpg	video	a video file format
mov	video	a video file format
avi	video	a video file format
hqx	archive	an archived file format
cpt	compression	a compressed file format

Content	Category	Description
application/octet-stream	unspecified data	byte stream
application/x-msdownload	Microsoft download	executable
application/exe	executable	executable
application/x-exe	executable	executable
application/dos-exe	DOS executable	executable
application/x-winexe	Windows executable	executable
application/msdos-windows	MS-DOS executable	executable
application/x-msdos-program	MS-DOS program	executable
application/x-oleobject	Microsoft OLE Object	executable
application/x-java-applet	Java Applet	executable
audio/mpegurl	MPEG audio URLs	audio
audio/x-mpegurl	MPEG audio URLs	audio
audio/mp3	MP3 audio	audio
audio/x-mp3	MP3 audio	audio
audio/mpeg	MPEG audio	audio
audio/mpg	MPEG audio	audio
audio/x-mpeg	MPEG audio	audio
audio/x-mpg	MPEG audio	audio
application/x-ogg	Ogg Vorbis	audio
audio/m4a	MPEG 4 audio	audio
audio/mp2	MP2 audio	audio

Content	Category	Description
audio/mp1	MP1 audio	audio
application/ogg	Ogg Vorbis	audio
audio/wav	Microsoft WAV	audio
audio/x-wav	Microsoft WAV	audio
audio/x-pn-wav	Microsoft WAV	audio
audio/aac	Advanced Audio Coding	audio
audio/midi	MIDI audio	audio
audio/mpeg	MPEG audio	audio
audio/aiff	AIFF audio	audio
audio/x-aiff	AIFF audio	audio
audio/x-pn-aiff	AIFF audio	audio
audio/x-pn-windows-acm	Windows ACM	audio
audio/x-pn-windows-pcm	Windows PCM	audio
audio/basic	8-bit u-law PCM	audio
audio/x-pn-au	Sun audio	audio
audio/3gpp	3GPP	audio
audio/3gpp-encrypted	encrypted 3GPP	audio
audio/scpls	streaming mp3 playlists	audio
audio/x-scpls	streaming mp3 playlists	audio
application/smil	SMIL	audio
application/sdp	Streaming Download Project	audio
application/x-sdp	Streaming Download Project	audio
audio/amr	AMR codec	audio
audio/amr-encrypted	AMR encrypted codec	audio
audio/amr-wb	AMR-WB codec	audio
audio/amr-wb-encrypted	AMR-WB encrypted codec	audio
audio/x-rn-3gpp-amr	3GPP codec	audio
audio/x-rn-3gpp-amr-encrypted	3GPP-AMR encrypted codec	audio
audio/x-rn-3gpp-amr-wb	3gpp-AMR-WB codec	audio
audio/x-rn-3gpp-amr-wb-encrypted	3gpp-AMR_WB encrypted codec	audio
application/streamingmedia	Streaming Media	audio
video/mpeg	MPEG video	video
audio/x-ms-wma	Windows Media	video

Content	Category	Description
video/quicktime	QuickTime	video
video/x-ms-asf	Microsoft ASF	video
video/x-msvideo	Microsoft AVI	video
video/x-sgi-mov	SGI movie	video
video/3gpp	3GPP video	video
video/3gpp-encrypted	3GPP encrypted video	video
video/3gpp2	3GPP2 video	video
audio/x-realaudio	RealAudio	audio
text/vnd.rn-realtext	RealText	text
audio/vnd.rn-realaudio	RealAudio	audio
audio/x-pn-realaudio	RealAudio plug-in	audio
image/vnd.rn-realpix	RealPix	image
application/vnd.rn-realmedia	RealMedia	video
application/vnd.rn-realmedia-vbr	RealMedia VBR	video
application/vnd.rn-realmedia-secure	secure RealMedia	video
application/vnd.rn-realaudio-secure	secure RealAudio	audio
audio/x-realaudio-secure	secure RealAudio	audio
video/vnd.rn-realvideo-secure	secure RealVideo	video
video/vnd.rn-realvideo	RealVideo	video
application/vnd.rn-realsystem-rmj	RealSystem media	video
application/vnd.rn-realsystem-rmx	RealSystem secure media	video
audio/rn-mpeg	MPEG audio	audio
application/x-shockwave-flash	Macromedia Shockwave	multimedia
application/x-director	Macromedia Shockwave	multimedia
application/x-authorware-bin	Macromedia Authorware binary	multimedia
application/x-authorware-map	Macromedia Authorware shocked file	multimedia
application/x-authorware-seg	Macromedia Authorware shocked packet	multimedia
application/futuresplash	Macromedia FutureSplash	multimedia
application/zip	ZIP	archive

Content	Category	Description
application/x-lzh	LZH archive	archive
image/gif	Graphics Interchange Format	image
image/png	Portable Network Graphics	image
image/jpeg	JPEG	image
image/bmp	Microsoft BMP	image
image/tiff	Tagged Image File Format	image
image/x-freehand	Macromedia Freehand	image
image/x-cmu-raster	CMU Raster	image
image/x-rgb	RGB image	image
text/css	cascading style sheet	text
text/html	HTML	text
text/plain	plain text	text
text/richtext	rich text	text
text/tab-separated-values	tab separated values	text
text/xml	XML	text
text/xsl	XSL	text
text/x-sgml	SGML	text
text/x-vcard	vCard	text
application/mac-binhex40	Macintosh BinHex	archive
application/x-stuffit	Macintosh Stuffit archive	archive
application/macwriteii	MacWrite Document	document
application/applefile	Macintosh File	archive
application/mac-compactpro	Macintosh Compact Pro	archive
application/x-bzip2	block compressed	compressed
application/x-shar	shell archive	archive
application/x-gtar	gzipped tar archive	archive
application/x-gzip	gzip compressed	compressed
application/x-tar	4.3BSD tar archive	archive
application/x-ustar	POSIX tar archive	archive
application/x-cpio	old cpio archive	archive
application/x-bcpio	POSIX cpio archive	archive
application/x-sv4crc	System V cpio with CRC	archive
application/x-compress	UNIX compressed	compressed
application/x-sv4cpio	System V cpio	archive

Content	Category	Description
application/x-sh	UNIX shell script	executable
application/x-csh	UNIX csh script	executable
application/x-tcl	Tcl script	executable
application/x-javascript	JavaScript	executable
application/x-excel	Microsoft Excel	document
application/mspowerpoint	Microsoft Powerpoint	document
application/msword	Microsoft Word	document
application/wordperfect5.1	Word Perfect	document
application/rtf	Rich Text Format	document
application/pdf	Adobe Acrobat	document
application/postscript	Postscript	document

Advanced

The Advanced section enables you to configure additional web filter options.

Secure Name Indication

- **Process HTTPS traffic by SNI (Server Name Indication) if present**: If this option is enabled, HTTPS traffic will be categorized using the "Server Name Indication" in the HTTPS data stream, if present. More details in HTTPS Options.
- **Process HTTPS traffic by hostname in server certificate when SNI information not present**: If this option is enabled *and* SNI information is not present, the certificate is fetched from the HTTPS server and the server name on the certificate will be used for categorization and filtering purposes.
- **Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available**: If this option is enabled *and* neither of the previous options worked, HTTPS traffic will be categorized using the IP address. More details in HTTPS Options.

Safe Browsing

- **Enforce safe search on popular search engines**: When this option is enabled, safe search will be enforced on all searches using supported search engines: Google, Yahoo!, Bing, Ask.
- **Enforce restrict mode on YouTube**: When this option is enabled, restrict mode will be enforced on all YouTube content.
- **Force searches through kid-friendly search engine**: When this option is enabled, all searches in popular search engines will be redirected through kidzsearch.com. kidzsearch is a visual child-safe search engine and web portal powered by Google Custom Search with academic autocomplete that emphasizes safety for children.

> **Note:** SSL Inspector must be installed and enabled to use all Safe browsing options.

Block Options

- **Block QUIC Sessions (UDP port 443)**: If enabled, Web Filter prevents browsers from using the QUIC protocol. QUIC is used by Chrome browser to access many types of Google applications and services. By allowing QUIC, NG Firewall has less visibility and control over this type of traffic.
- **Block pages from IP only hosts**: When this option is enabled, users entering an IP address rather than domain name will be blocked.
- **Pass if referrer matches any Pass Sites**. When this option is checked, if a site allowed via a Pass Site entry links to external content using a referrer, that external content will be passed regardless

of Category settings. Please note that this setting requires SSL Inspector to work with HTTPS sites and many sites no longer use referrers.

- **Close connection for blocked HTTPS sessions without redirecting to block page**. If enabled, secure sites blocked by Web Filter do not redirect the user to a denial page and close the connection without any notice to the user. This is useful when you are not using SSL Inspector and the server's root certificate authority is not installed on the client device.

Google Restrictions

- **Restrict Google applications**: When this option is enabled, only domains listed in **Domain** are allowed to access Google applications such as Gmail. All others are blocked by Google. Multiple domains can be specified, separated by commas such as **google.com,domain.com**. This adds an X-GoogApps-Allowed-Domains header to web requests which then gets enforced on Google's servers. More information on this feature can be found here.

> **Note:** SSL Inspector must be installed and enabled to restrict Google applications.

Custom Block Page

- **Custom block page URL**: Set an external location where to redirect users when they are denied access to a web site by Web Filter. This is useful if you want to use your own server to process the denial in a different way than the built-in denial options.

The following query string variables are passed to the forward location so the receiving system can process the information.

Variable	Description
reason	The reason the user was denied access.
appname	The NG Firewall app that is responsible for the denial.
appid	The ID of the NG Firewall app that is responsible for the denial.
clientAddress	The IP Address of the denied device.
url	The denied URL that the user requested.

Unblock Options

- **Unblock**: This section can be used to add a button to allow users to bypass restrictions on a case-by-case basis.

    If Unblock is set to **None**, no users will be allowed to bypass the block page. If Unblock is set to **Temporary**, users will be allowed to visit the site for one hour from the time it is unblocked. If Unblock is set to **Permanent and Global**, then users will be allowed to visit the site and unblocked sites will be added to the permanent global pass list so it will always be allowed in the future.

    You also have the option of setting a password to Unblock; it can either be the existing Administrator password for the NG Firewall or you can set a new, separate password only for the Unblock feature.

- **Clear Category URL Cache**: This option will clear the local cache of categorized sites and URLs. After clearing the cache all new web visits will be looked up fresh using the categorization service. The cache automatically cleans itself as entries become old or stale, so this is mostly for testing.
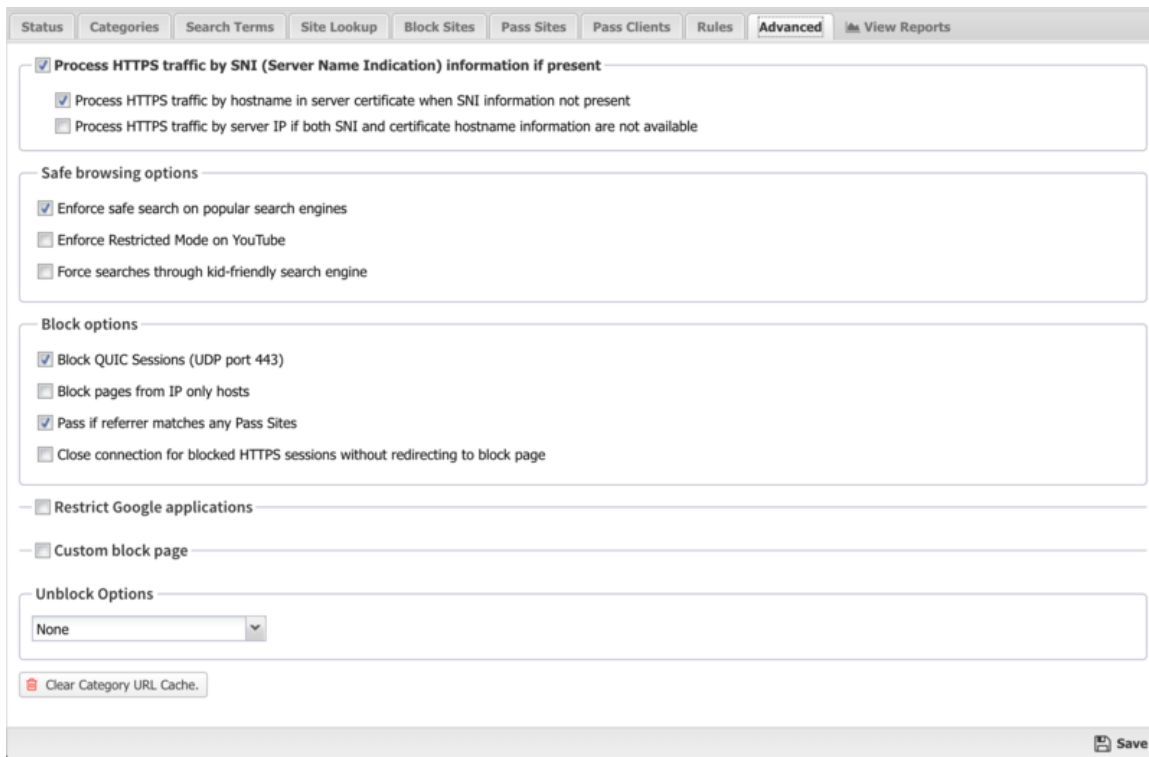
Figure 69: Web Filter Advanced tab

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Web Filter.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Web Filter Summary	A summary of web filter actions.
Web Usage	The amount of total, flagged, and blocked web requests over time.
Web Usage (scanned)	The amount of total, flagged, and blocked web requests over time.
Web Usage (flagged)	The amount of flagged, and blocked web requests over time.
Web Usage (blocked)	The amount of flagged, and blocked web requests over time.
Top Categories (by request)	The number of web requests grouped by category.
Top Categories (by size)	The sum of the size of requested web content grouped by category.
Top Flagged Categories	The number of flagged web requests grouped by category.
Top Blocked Categories	The number of blocked web requests grouped by category.

Report Entry	Description
Top Sites (by request)	The number of web requests grouped by website.
Top Sites (by size)	The sum of the size of requested web content grouped by website.
Top Flagged Sites	The number of flagged web requests grouped by website.
Top Blocked Sites	The number of blocked web requests grouped by website.
Top Domains (by request)	The number of web requests grouped by domain.
Top Domains (by size)	The sum of the size of requested web content grouped by domain.
Top Flagged Domains	The number of flagged web requests grouped by domain.
Top Blocked Domains	The number of blocked web requests grouped by domain.
Top Domains Usage	The amount of web requests per top domain.
Top Hostnames (by requests)	The number of web requests grouped by hostname.
Top Hostnames (by size)	The sum of the size of requested web content grouped by hostname.
Top Flagged Hostnames	The number of flagged web request grouped by hostname.
Top Blocked Hostnames	The number of blocked web request grouped by hostname.
Top Clients (by requests)	The number of web requests grouped by client.
Top Clients (by size)	The sum of the size of requested web content grouped by client.
Top Flagged Clients	The number of flagged web request grouped by client.
Top Blocked Clients	The number of blocked web request grouped by client.
Top Usernames (by requests)	The number of web requests grouped by username.
Top Usernames (by size)	The sum of the size of requested web content grouped by username.
Top Flagged Usernames	The number of flagged web request grouped by username.
Top Blocked Usernames	The number of blocked web request grouped by username.
Top Content (by request)	The number of web requests grouped by category.
Top Content (by size)	The sum of the size of requested web content grouped by category.
Top Searches	The number of non blocked, non-flagged search queries grouped by term.
Top Flagged Searches	The number of flagged search queries grouped by term.
Top Blocked Searches	The number of blocked search queries grouped by term.
All Web Events	Shows all scanned web requests.
Flagged Web Events	Shows all flagged web requests.
Blocked Web Events	Shows all blocked web requests.
All HTTP Events	Shows all scanned unencrypted HTTP requests.

Report Entry	Description
All HTTPS Events	Shows all encrypted HTTPS requests.
Unblocked Web Events	Shows all unblocked web requests
All Search Events	Shows all search queries processed by Web Filter.
Flagged Search Events	Shows flagged search queries processed by Web Filter.
Blocked Search Events	Shows blocked search queries processed by Web Filter.

The tables queried to render these reports:

- Sessions
- Session Minutes

HTTPS Options

There are many ways to handle HTTPS. An overview of the various techniques is described here.

If SSL Inspector is installed and inspects a session, then it is fully decrypted to HTTP before Web Filter processes the session. In this case HTTPS is treated identically to HTTP. If SSL Inspector is not installed or the session is not inspected, there are still several techniques to handle encrypted HTTP sessions.

There are three HTTPS options.

- Process HTTPS traffic by SNI (Server Name Indication) if present.
- Process HTTPS traffic by hostname in server certificate when SNI information not present
- Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available.

If *Process HTTPS traffic by SNI (Server Name Indication) if present* encrypted **port-443** traffic will be scanned. Most modern browsers on modern OSs will send the hostname of the server in cleartext - this is called "Server Name Indication" or SNI. SNI is an optional cleartext field in the HTTPS request that shows the hostname of the server. If this option is enabled and the SNI information is present in the HTTPS request, this hostname will be used as the URL for this request and all categorization, flag lists, and pass lists, will be processed as if this were a regular HTTP request to that URL.

If the SNI-based categorization determines the page should be passed (and/or flagged) then the session is allowed and the appropriate event based on the SNI information is logged ("https://example.com/").

> For example, if the user visits "https://wellsfargo.com/welcome" in the browser, "wellsfargo.com" is seen as the SNI information. If SNI-based categorization is enabled, the request will be handled exactly like "http://wellsfargo.com" would be. If Web Filter is configured to flag "Financial Services" , then "https://wellsfargo.com/welcome" will be flagged, unless "wellsfargo.com" is in the pass list or the client IP address is in the client IP pass list.

If No SNI information is present and *Process HTTPS traffic by hostname in server certificate when SNI information not present* is enabled, then the hostname will be pulled from the certificate presented to the client.

> For example, if the user visits "https://wellsfargo.com/welcome" in a non-SNI enabled browser, then there is no SNI information. In this case if *Process HTTPS traffic by hostname in server certificate when SNI information not present* is enabled it will use the certificate information instead to categorize the session. It will download the certificate from the site and see that the certificate is "Issued To" "www.wellsfargo.com." It will use this information to check the category for "https://www.wellsfargo.com" and categorize the session.

If no SNI or certificate information was available and *Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available* the session will be processed and categorized by IP address. If the IP-based processing and categorization of the web requests determines the session should be flagged, the session is reset and no more processing of this session will be done. If the IP-based processing and categorization determines the page should be passed (and/or flagged) then the session is allowed and the appropriate event based on its IP is logged ("https://1.2.3.4").

> For example, if the user visits "https://wellsfargo.com/welcome" in a non-SNI enabled browser, then there is no SNI information. If the the certificate information was missing for some reason then this session can only be identified by IP address. In this case if *Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available* is enabled it will use the IP address instead. So it will process/categorize this web request as 'http://1.2.3.4' if *1.2.3.4* is the IP of wellsfargo.com. This will still often result in correct categorization for dedicated web servers, but does poorly when using generic cloud computing servers that offer a wide variety of websites.

> **Note:** Neither HTTPS processing (SNI, certificate, or IP-based categorization) can read the URI information as it is not sent in cleartext. As such the URI will not be used as part of the categorization and the URI is assumed to be "/" when evaluating pass rules. If scanning the URI is necessary then full SSL Inspection may be required. Read HTTPS.

To see the HTTPS categorization in action use the "All HTTPS Events" query in the event log.

8.7.1   Web Filter Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Web Filter.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Web Filter Summary	A summary of web filter actions.
Web Usage	The amount of total, flagged, and blocked web requests over time.
Web Usage (scanned)	The amount of total, flagged, and blocked web requests over time.
Web Usage (flagged)	The amount of flagged, and blocked web requests over time.
Web Usage (blocked)	The amount of flagged, and blocked web requests over time.
Top Categories (by request)	The number of web requests grouped by category.
Top Categories (by size)	The sum of the size of requested web content grouped by category.
Top Flagged Categories	The number of flagged web requests grouped by category.
Top Blocked Categories	The number of blocked web requests grouped by category.
Top Sites (by request)	The number of web requests grouped by website.

Report Entry	Description
Top Sites (by size)	The sum of the size of requested web content grouped by website.
Top Flagged Sites	The number of flagged web requests grouped by website.
Top Blocked Sites	The number of blocked web requests grouped by website.
Top Domains (by request)	The number of web requests grouped by domain.
Top Domains (by size)	The sum of the size of requested web content grouped by domain.
Top Flagged Domains	The number of flagged web requests grouped by domain.
Top Blocked Domains	The number of blocked web requests grouped by domain.
Top Domains Usage	The amount of web requests per top domain.
Top Hostnames (by requests)	The number of web requests grouped by hostname.
Top Hostnames (by size)	The sum of the size of requested web content grouped by hostname.
Top Flagged Hostnames	The number of flagged web request grouped by hostname.
Top Blocked Hostnames	The number of blocked web request grouped by hostname.
Top Clients (by requests)	The number of web requests grouped by client.
Top Clients (by size)	The sum of the size of requested web content grouped by client.
Top Flagged Clients	The number of flagged web request grouped by client.
Top Blocked Clients	The number of blocked web request grouped by client.
Top Usernames (by requests)	The number of web requests grouped by username.
Top Usernames (by size)	The sum of the size of requested web content grouped by username.
Top Flagged Usernames	The number of flagged web request grouped by username.
Top Blocked Usernames	The number of blocked web request grouped by username.
Top Content (by request)	The number of web requests grouped by category.
Top Content (by size)	The sum of the size of requested web content grouped by category.
Top Searches	The number of non blocked, non-flagged search queries grouped by term.
Top Flagged Searches	The number of flagged search queries grouped by term.
Top Blocked Searches	The number of blocked search queries grouped by term.
All Web Events	Shows all scanned web requests.
Flagged Web Events	Shows all flagged web requests.
Blocked Web Events	Shows all blocked web requests.
All HTTP Events	Shows all scanned unencrypted HTTP requests.
All HTTPS Events	Shows all encrypted HTTPS requests.

Report Entry	Description
Unblocked Web Events	Shows all unblocked web requests
All Search Events	Shows all search queries processed by Web Filter.
Flagged Search Events	Shows flagged search queries processed by Web Filter.
Blocked Search Events	Shows blocked search queries processed by Web Filter.

The tables queried to render these reports:

- Sessions
- Session Minutes

Related Topics

- Report Viewer
- Reports

8.8 Web Monitor

Web Monitor monitors HTTP and HTTPS traffic on your network to log web activities and flag inappropriate content.

About Web Monitor

- **Real-time classification and updates**: When your users visit a site, NG Firewall sends the URL to the Webroot BrightCloud® to be categorized. When the data is returned, NG Firewall keeps a temporary local cache of the site and category to speed up the process the next time the URL is requested. This data is then used to flag or allow users access to the site they have requested, all without any appreciable increase in load time. If a site is not categorized upon request, it is autocategorized by our partners at Webroot and put into a queue to be verified by a human. Because this is done dynamically, new sites and updated URLs are allowed or flagged according to your settings without additional intervention, plus you have the option of requesting recategorization of sites.
- **HTTPS Filtering**: Web Monitor has multiple techniques to deal with HTTPS, SSL-encrypted HTTP. HTTPS traffic is encrypted so only some information is visible and this information is used to categorize the session. More information on how this is down below.
- **Detailed categorization**: Web Monitor offers 79 categories and tens of billions of URLs. The Web Monitor database is over 100 times larger and more accurate. The abundance of categories means that you can narrow your scope - maybe you want to flag websites related to nudity, but allow sites dealing with Sexual Education.

Traffic Flow

When scanning traffic, Web Monitor evaluates the pass lists, flag lists, categories, and rules at two distinct points of the HTTP transaction. The first evaluation happens after the request is received from the client and before it is forwarded to the server. The second is after the response is received from the server and before it is passed back to the client. This allows a high degree of monitoring over both resources that are requested, and content that is returned in response.

HTTP Request

When evaluating HTTP requests, Web Monitor applies the configured rules and lists in the following order:

1. A lookup is performed to determine the category for the requested site. The category is attached to the session for use by Web Monitor as well as other applications.
2. The source IP of the request is checked against the Pass Clients list. If a match is found, the traffic is allowed.
3. The destination site of the request is checked against the Pass Sites list. If a match is found, the traffic is allowed.
4. The destination site of the request is checked against the Flag Sites list. If a match is found, the traffic is flagged.
5. The traffic details are passed to the Rules list. If a match is found, the traffic is allowed and possibly flagged based on the options configured in the rule that was matched.
6. The category determined in step #1 is compared to the Categories list, and the traffic is allowed and possibly flagged based on the corresponding match. If the category could not be determined, the traffic is allowed.

HTTP Response

When evaluating HTTP responses, Web Monitor applies the configured rules and lists in the following order:

1. The source IP of the request is checked against the Pass Clients list. If a match is found, the traffic is allowed.
2. The site from which the response was received is checked against the Pass Sites list. If a match is found, the traffic is allowed.
3. The traffic details are passed to the Rules list. If a match is found, the traffic is allowed and possibly flagged based on the options configured in the rule that was matched.

Settings

This section reviews the different settings and configuration options available for Web Monitor.

Status

This displays the current status and some statistics.

Figure 70: Web Monitor Status

Categories

Categories allows you to customize which categories of sites will be flagged. Categories that are flagged will allow the user to access the site, but will be silently flagged as a violation for event logs and Reports. These flag actions operate the same way for all of the different Web Monitor options.
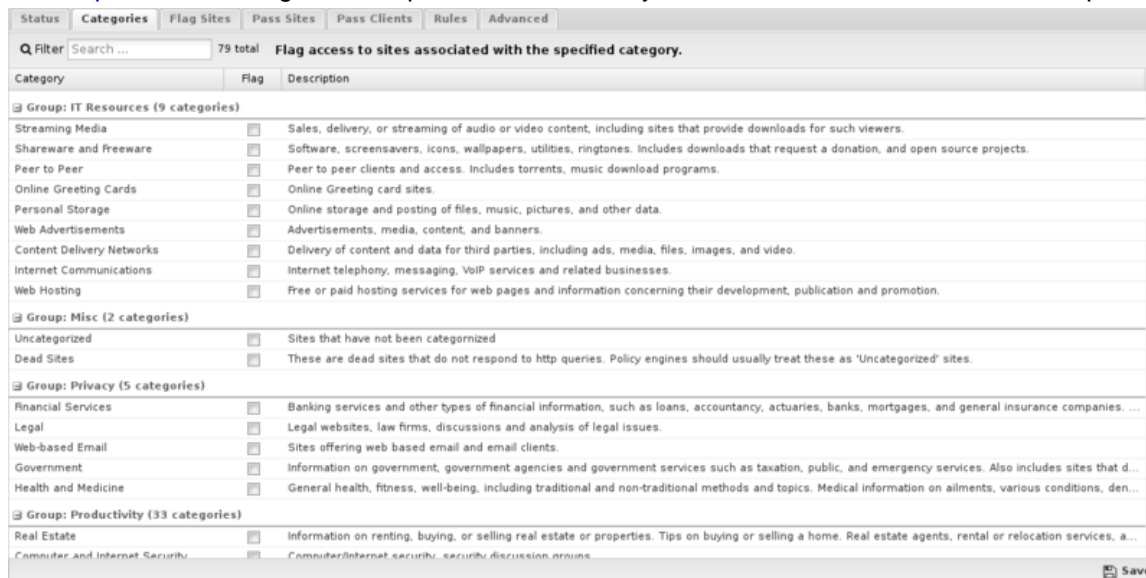
Figure 71: Web Monitor Categories

Flag Sites

Under Flag Sites you can add individual domain names you want to be flagged - just enter the domain name (e.g. youtube.com) and specify your chosen action. This list uses URL Matcher syntax.
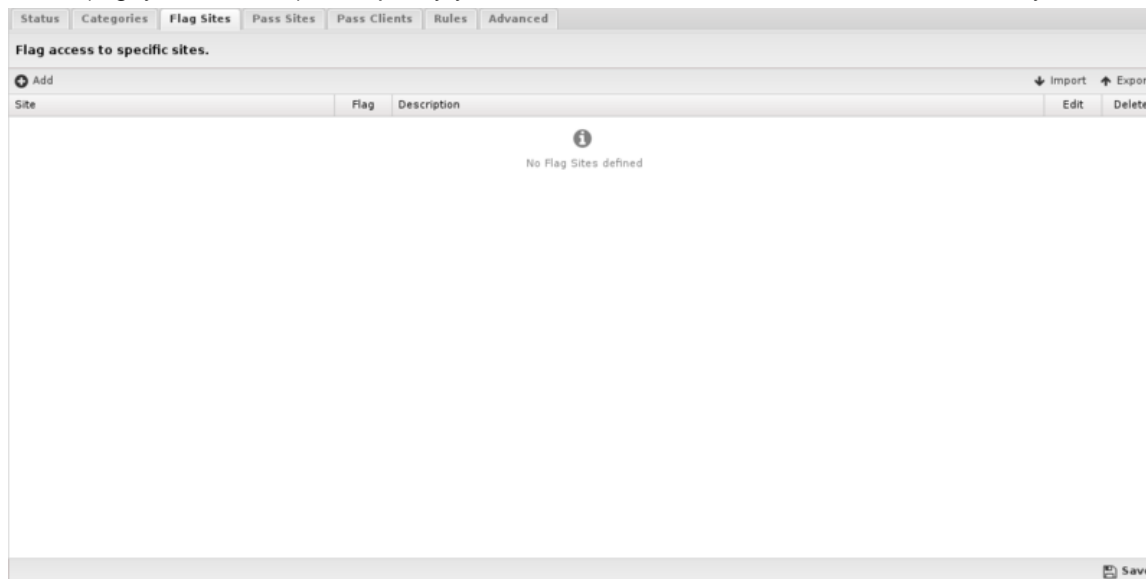
Figure 72: Web Monitor Flag Sites

Pass Sites

Pass Sites is used to pass content that would have otherwise been flagged. This can be useful for "unflagging" sites that you don't want flagged according to flag settings. Any domains you add to the Passed Sites list will be allowed, even if flagged by category or by individual URL - just add the domain

and save. Unchecking the pass option will allow the site to be flagged as if the entry was not present. This list uses URL Matcher syntax.
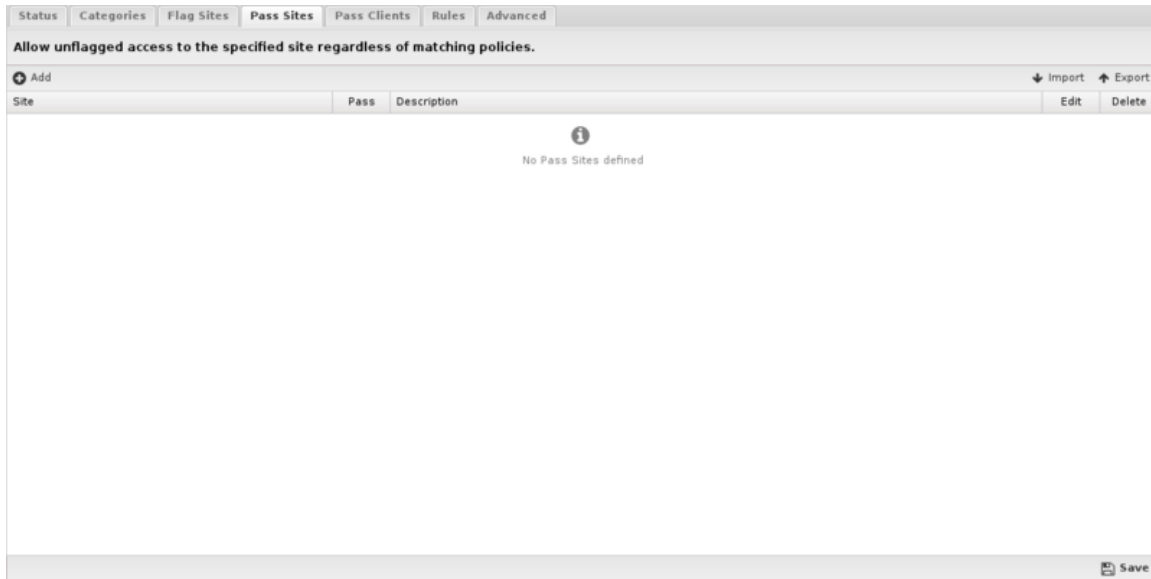
Figure 73: Web Monitor Pass Sites

Pass Clients

If you add an IP address to this list, Web Monitor will not flag any traffic from that IP regardless of the flagged categories or sites. Just add the IP and save. Unchecking the pass option will have the flag/pass lists affect the user as if they were not entered into the Passed Client IPs list. This list uses IP Matcher syntax.

> If you have a few users that need to completely bypass Web Monitor controls, consider using pass lists. If you have users that simply need different Web Monitor settings, you should set up a separate rack using Policy Manager. When using this feature, please remember that DHCP IPs can change, so you'll probably want to set up either a Static IP or a Static DHCP Lease for the machine in question.

Figure 74: Web Monitor Pass Clients

Rules

The **Rules** tab allows you to specify rules to Flag traffic that passes through Web Monitor.

The Rules describes how rules work and how they are configured. Web Monitor uses rules to determine when to flag specific sessions. Flagging a session marks it in the logs for reviewing in the event logs or reports, but has no direct effect on the network traffic.
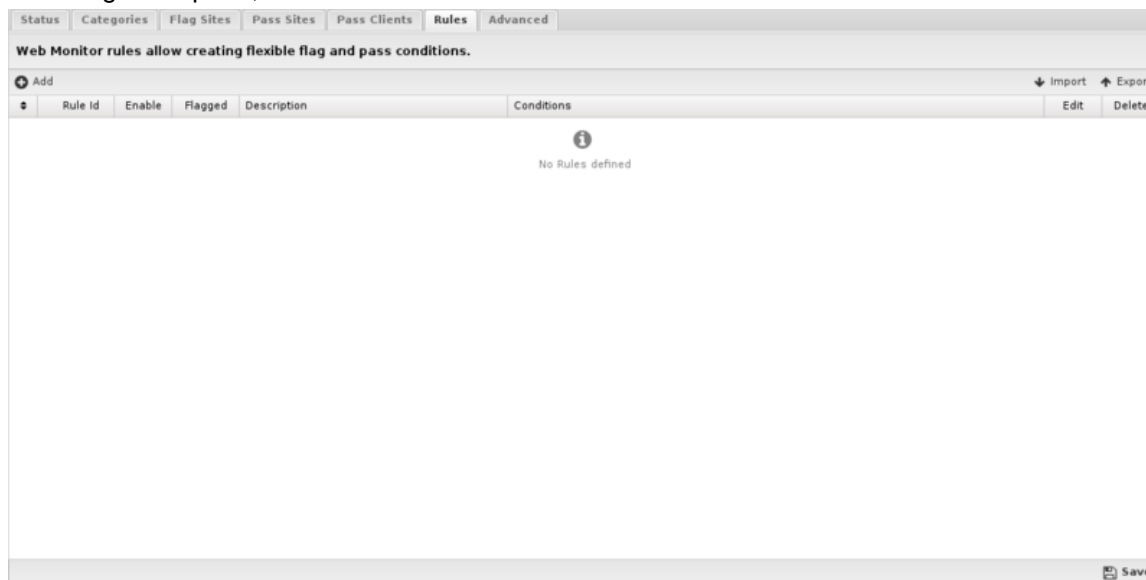
Figure 75: Web Monitor Rules

Rule Actions

- **Flag**: Allows the traffic which matched the rule to flow, and flags the traffic for easier viewing in the event log.

Rule Types

In previous versions of Web Monitor, there were dedicated lists for flagging certain file extensions or MIME types. This capability is still available using the more flexible filter rules. For flagging specific file extensions, you can create a rule with the condition **Web Filter: Response File Extension** that has a comma separated list of the extensions to flag in the Value field. For flagging MIME types, you would create a rule with the condition **Web Filter: Response Content Type** that has a comma separated list of the content types to flag in the Value field.

Below are tables that list the default file extensions and MIME types that were available in previous versions. Note that these lists are not exhaustive, but are included here as a reference, and to simplify creation of such rules via copy/paste of the values in the tables.

Extension	Category	Description
exe	executable	an executable file format
ocx	executable	an executable file format
dll	executable	an executable file format
cab	executable	an ActiveX executable file format
bin	executable	an executable file format
com	executable	an executable file format
jpg	image	an image file format

Extension	Category	Description
png	image	an image file format
gif	image	an image file format
jar	java	a Java file format
class	java	a Java file format
swf	flash	the flash file format
mp3	audio	an audio file format
wav	audio	an audio file format
wmf	audio	an audio file format
mpg	video	a video file format
mov	video	a video file format
avi	video	a video file format
hqx	archive	an archived file format
cpt	compression	a compressed file format

Content	Category	Description
application/octet-stream	unspecified data	byte stream
application/x-msdownload	Microsoft download	executable
application/exe	executable	executable
application/x-exe	executable	executable
application/dos-exe	DOS executable	executable
application/x-winexe	Windows executable	executable
application/msdos-windows	MS-DOS executable	executable
application/x-msdos-program	MS-DOS program	executable
application/x-oleobject	Microsoft OLE Object	executable
application/x-java-applet	Java Applet	executable
audio/mpegurl	MPEG audio URLs	audio
audio/x-mpegurl	MPEG audio URLs	audio
audio/mp3	MP3 audio	audio
audio/x-mp3	MP3 audio	audio
audio/mpeg	MPEG audio	audio
audio/mpg	MPEG audio	audio
audio/x-mpeg	MPEG audio	audio
audio/x-mpg	MPEG audio	audio
application/x-ogg	Ogg Vorbis	audio

Content	Category	Description
audio/m4a	MPEG 4 audio	audio
audio/mp2	MP2 audio	audio
audio/mp1	MP1 audio	audio
application/ogg	Ogg Vorbis	audio
audio/wav	Microsoft WAV	audio
audio/x-wav	Microsoft WAV	audio
audio/x-pn-wav	Microsoft WAV	audio
audio/aac	Advanced Audio Coding	audio
audio/midi	MIDI audio	audio
audio/mpeg	MPEG audio	audio
audio/aiff	AIFF audio	audio
audio/x-aiff	AIFF audio	audio
audio/x-pn-aiff	AIFF audio	audio
audio/x-pn-windows-acm	Windows ACM	audio
audio/x-pn-windows-pcm	Windows PCM	audio
audio/basic	8-bit u-law PCM	audio
audio/x-pn-au	Sun audio	audio
audio/3gpp	3GPP	audio
audio/3gpp-encrypted	encrypted 3GPP	audio
audio/scpls	streaming mp3 playlists	audio
audio/x-scpls	streaming mp3 playlists	audio
application/smil	SMIL	audio
application/sdp	Streaming Download Project	audio
application/x-sdp	Streaming Download Project	audio
audio/amr	AMR codec	audio
audio/amr-encrypted	AMR encrypted codec	audio
audio/amr-wb	AMR-WB codec	audio
audio/amr-wb-encrypted	AMR-WB encrypted codec	audio
audio/x-rn-3gpp-amr	3GPP codec	audio
audio/x-rn-3gpp-amr-encrypted	3GPP-AMR encrypted codec	audio
audio/x-rn-3gpp-amr-wb	3gpp-AMR-WB codec	audio
audio/x-rn-3gpp-amr-wb-encrypted	3gpp-AMR_WB encrypted codec	audio
application/streamingmedia	Streaming Media	audio
video/mpeg	MPEG video	video

Content	Category	Description
audio/x-ms-wma	Windows Media	video
video/quicktime	QuickTime	video
video/x-ms-asf	Microsoft ASF	video
video/x-msvideo	Microsoft AVI	video
video/x-sgi-mov	SGI movie	video
video/3gpp	3GPP video	video
video/3gpp-encrypted	3GPP encrypted video	video
video/3gpp2	3GPP2 video	video
audio/x-realaudio	RealAudio	audio
text/vnd.rn-realtext	RealText	text
audio/vnd.rn-realaudio	RealAudio	audio
audio/x-pn-realaudio	RealAudio plug-in	audio
image/vnd.rn-realpix	RealPix	image
application/vnd.rn-realmedia	RealMedia	video
application/vnd.rn-realmedia-vbr	RealMedia VBR	video
application/vnd.rn-realmedia-secure	secure RealMedia	video
application/vnd.rn-realaudio-secure	secure RealAudio	audio
audio/x-realaudio-secure	secure RealAudio	audio
video/vnd.rn-realvideo-secure	secure RealVideo	video
video/vnd.rn-realvideo	RealVideo	video
application/vnd.rn-realsystem-rmj	RealSystem media	video
application/vnd.rn-realsystem-rmx	RealSystem secure media	video
audio/rn-mpeg	MPEG audio	audio
application/x-shockwave-flash	Macromedia Shockwave	multimedia
application/x-director	Macromedia Shockwave	multimedia
application/x-authorware-bin	Macromedia Authorware binary	multimedia
application/x-authorware-map	Macromedia Authorware shocked file	multimedia
application/x-authorware-seg	Macromedia Authorware shocked packet	multimedia
application/futuresplash	Macromedia FutureSplash	multimedia
application/zip	ZIP	archive
application/x-lzh	LZH archive	archive
image/gif	Graphics Interchange Format	image

Content	Category	Description
image/png	Portable Network Graphics	image
image/jpeg	JPEG	image
image/bmp	Microsoft BMP	image
image/tiff	Tagged Image File Format	image
image/x-freehand	Macromedia Freehand	image
image/x-cmu-raster	CMU Raster	image
image/x-rgb	RGB image	image
text/css	cascading style sheet	text
text/html	HTML	text
text/plain	plain text	text
text/richtext	rich text	text
text/tab-separated-values	tab separated values	text
text/xml	XML	text
text/xsl	XSL	text
text/x-sgml	SGML	text
text/x-vcard	vCard	text
application/mac-binhex40	Macintosh BinHex	archive
application/x-stuffit	Macintosh Stuffit archive	archive
application/macwriteii	MacWrite Document	document
application/applefile	Macintosh File	archive
application/mac-compactpro	Macintosh Compact Pro	archive
application/x-bzip2	block compressed	compressed
application/x-shar	shell archive	archive
application/x-gtar	gzipped tar archive	archive
application/x-gzip	gzip compressed	compressed
application/x-tar	4.3BSD tar archive	archive
application/x-ustar	POSIX tar archive	archive
application/x-cpio	old cpio archive	archive
application/x-bcpio	POSIX cpio archive	archive
application/x-sv4crc	System V cpio with CRC	archive
application/x-compress	UNIX compressed	compressed
application/x-sv4cpio	System V cpio	archive
application/x-sh	UNIX shell script	executable
application/x-csh	UNIX csh script	executable

Content	Category	Description
application/x-tcl	Tcl script	executable
application/x-javascript	JavaScript	executable
application/x-excel	Microsoft Excel	document
application/mspowerpoint	Microsoft Powerpoint	document
application/msword	Microsoft Word	document
application/wordperfect5.1	Word Perfect	document
application/rtf	Rich Text Format	document
application/pdf	Adobe Acrobat	document
application/postscript	Postscript	documen

Advanced

The Advanced section allows you to configure additional Web Monitor options.

- **Process HTTPS traffic by SNI (Server Name Indication) if present**: If this option is enabled, HTTPS traffic will be categorized using the "Server Name Indication" in the HTTPS data stream, if present. More details in HTTPS Options.
- **Process HTTPS traffic by hostname in server certificate when SNI information not present**: If this option is enabled *and* SNI information is not present, the certificate is fetched from the HTTPS server and the server name on the certificate will be used for categorization and filtering purposes.
- **Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available**: If this option is enabled *and* neither of the previous options worked, HTTPS traffic will be categorized using the IP address. More details in HTTPS Options.
- **Clear Category URL Cache**: This option will clear the local cache of categorized sites and URLs. After clearing the cache all new web visits will be looked up fresh using the categorization service. The cache automatically cleans itself as entries become old or stale, so this is mostly for testing.

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Web Monitor.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Web Monitor Summary	A summary of web monitor actions.
Web Usage	The amount of total and flagged web requests over time.
Web Usage (scanned)	The amount of total web requests over time.
Web Usage (flagged)	The amount of flagged web requests over time.
Top Categories (by request)	The number of web requests grouped by category.
Top Categories (by size)	The sum of the size of requested web content grouped by category.
Top Flagged Categories	The number of flagged web requests grouped by category.
Top Sites (by request)	The number of web requests grouped by website.
Top Sites (by size)	The sum of the size of requested web content grouped by website.
Top Flagged Sites	The number of flagged web requests grouped by website.
Top Domains (by request)	The number of web requests grouped by domain.
Top Domains (by size)	The sum of the size of requested web content grouped by domain.
Top Flagged Domains	The number of flagged web requests grouped by domain.
Top Domains Usage	The amount of web requests per top domain.
Top Hostnames (by requests)	The number of web requests grouped by hostname.
Top Hostnames (by size)	The sum of the size of requested web content grouped by hostname.
Top Flagged Hostnames	The number of flagged web request grouped by hostname.
Top Clients (by requests)	The number of web requests grouped by client.
Top Clients (by size)	The sum of the size of requested web content grouped by client.
Top Flagged Clients	The number of flagged web request grouped by client.
Top Usernames (by requests)	The number of web requests grouped by username.
Top Usernames (by size)	The sum of the size of requested web content grouped by username.
Top Flagged Usernames	The number of flagged web request grouped by username.
Top Content (by request)	The number of web requests grouped by category.
Top Content (by size)	The sum of the size of requested web content grouped by category.

Report Entry	Description
Top Searches	The number of non blocked, non-flagged search queries grouped by term.
Top Flagged Searches	The number of flagged search queries grouped by term.
All Web Events	Shows all scanned web requests.
Flagged Web Events	Shows all flagged web requests.
All HTTP Events	Shows all scanned unencrypted HTTP requests.
All HTTPS Events	Shows all encrypted HTTPS requests.
All Search Events	Shows all search queries processed by Web Monitor.
Flagged Search Events	Shows flagged search queries processed by Web Monitor.

The tables queried to render these reports:

- Sessions
- Session Minutes

HTTPS Options

There are many ways to handle HTTPS. An overview of the various techniques is described here.

If SSL Inspector is installed and inspects a session, then it is fully decrypted to HTTP before Web Monitor processes the session. In this case HTTPS is treated identically to HTTP. If SSL Inspector is not installed or the session is not inspected, there are still several techniques to handle encrypted HTTP sessions.

There are three HTTPS options.

- Process HTTPS traffic by SNI (Server Name Indication) if present.
- Process HTTPS traffic by hostname in server certificate when SNI information not present
- Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available.

If *Process HTTPS traffic by SNI (Server Name Indication) if present* encrypted port-443 traffic will be scanned. Most modern browsers on modern OSs will send the hostname of the server in cleartext - this is called "Server Name Indication" or SNI. SNI is an optional cleartext field in the HTTPS request that shows the hostname of the server. If this option is enabled and the SNI information is present in the HTTPS request, this hostname will be used as the URL for this request and all categorization, flag lists, and pass lists, will be processed as if this were a regular HTTP request to that URL.

If the SNI-based categorization determines the page should be passed (and/or flagged) then the session is allowed and the appropriate event based on the SNI information is logged ("https://example.com/").

> For example, if the user visits "https://wellsfargo.com/welcome" in the browser, it will see "wellsfargo.com" as the SNI information. If enabled, the request will be handled exactly like "http://wellsfargo.com" would be. If "Banking" is flagged it will be flagged, unless "wellsfargo.com" is in the pass list or the client IP is in the client IP pass list. If "wellsfargo.com" is flaggeed it will be flagged, unless "wellsfargo.com" is in the pass list or the client IP is in the client IP pass list.

If No SNI information is present and *Process HTTPS traffic by hostname in server certificate when SNI information not present* is enabled, then the hostname will be pulled from the certificate presented to the client.

For example, if the user visits "https://wellsfargo.com/welcome" in a non-SNI enabled browser, then there is no SNI information. In this case if *Process HTTPS traffic by hostname in server certificate when SNI information not present* is enabled it will use the certificate information instead to categorize the session. It will download the certificate from the site and see that the certificate is "Issued To" "www.wellsfargo.com." It will use this information to check the category for "https://www.wellsfargo.com" and categorize the session.

If no SNI or certificate information was available and *Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available* the session will be processed and categorized by IP address. If the IP-based processing and categorization of the web requests determines the session should be flagged, the session is reset and no more processing of this session will be done. If the IP-based processing and categorization determines the page should be passed (and/or flagged) then the session is allowed and the appropriate event based on its IP is logged ("https://1.2.3.4").

For example, if the user visits "https://wellsfargo.com/welcome" in a non-SNI enabled browser, then there is no SNI information. If the the certificate information was missing for some reason then this session can only be identified by IP address. In this case if *Process HTTPS traffic by server IP if both SNI and certificate hostname information are not available* is enabled it will use the IP address instead. So it will process/categorize this web request as 'http://1.2.3.4' if *1.2.3.4* is the IP of wellsfargo.com. This will still often result in correct categorization for dedicated web servers, but does poorly when using generic cloud computing servers that offer a wide variety of websites.

> **Note:** Neither HTTPS process (SNI, certificate, or IP-based categorization) can read the URI information as it is not sent in cleartext. As such the URI will not be used as part of the categorization and the URI is assumed to be "/" when evaluating pass rules. If scanning the URI is necessary then full SSL Inspection may be required.

To see the HTTPS categorization in action use the "All HTTPS Events" query in the event log.

8.8.1   Web Monitor Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Web Monitor.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Web Monitor Summary	A summary of web monitor actions.
Web Usage	The amount of total and flagged web requests over time.
Web Usage (scanned)	The amount of total web requests over time.
Web Usage (flagged)	The amount of flagged web requests over time.
Top Categories (by request)	The number of web requests grouped by category.
Top Categories (by size)	The sum of the size of requested web content grouped by category.
Top Flagged Categories	The number of flagged web requests grouped by category.
Top Sites (by request)	The number of web requests grouped by website.

Report Entry	Description
Top Sites (by size)	The sum of the size of requested web content grouped by website.
Top Flagged Sites	The number of flagged web requests grouped by website.
Top Domains (by request)	The number of web requests grouped by domain.
Top Domains (by size)	The sum of the size of requested web content grouped by domain.
Top Flagged Domains	The number of flagged web requests grouped by domain.
Top Domains Usage	The amount of web requests per top domain.
Top Hostnames (by requests)	The number of web requests grouped by hostname.
Top Hostnames (by size)	The sum of the size of requested web content grouped by hostname.
Top Flagged Hostnames	The number of flagged web request grouped by hostname.
Top Clients (by requests)	The number of web requests grouped by client.
Top Clients (by size)	The sum of the size of requested web content grouped by client.
Top Flagged Clients	The number of flagged web request grouped by client.
Top Usernames (by requests)	The number of web requests grouped by username.
Top Usernames (by size)	The sum of the size of requested web content grouped by username.
Top Flagged Usernames	The number of flagged web request grouped by username.
Top Content (by request)	The number of web requests grouped by category.
Top Content (by size)	The sum of the size of requested web content grouped by category.
Top Searches	The number of non blocked, non-flagged search queries grouped by term.
Top Flagged Searches	The number of flagged search queries grouped by term.
All Web Events	Shows all scanned web requests.
Flagged Web Events	Shows all flagged web requests.
All HTTP Events	Shows all scanned unencrypted HTTP requests.
All HTTPS Events	Shows all encrypted HTTPS requests.
All Search Events	Shows all search queries processed by Web Monitor.
Flagged Search Events	Shows flagged search queries processed by Web Monitor.

The tables queried to render these reports:

- Sessions
- Session Minutes

Related Topics

Report Viewer

Reports

NG Firewall Protect Apps

Contents

9.1   Firewall

Firewall provides traditional firewall functionality, blocking and/or flagging traffic based on rules.

The term "Firewall" has grown to encompass many functionalities and has a wide array of meanings. The "firewall" is often use interchangeably with **router**, **gateway** and **Unified Threat Management (UTM)** or . Even the NG Firewall is a "next-gen" "firewall." There are also host-based "firewalls" that run on the local host computer.

The "Firewall" app itself is a traditional firewall used to block and/or flag **TCP** and **UDP** sessions passing through NG Firewall using rules. The Firewall app provides the same functionality as the traditional "firewall" - the ability to use rules to control which computers and communicate on a network.

Settings

This section reviews the different settings and configuration options available for Firewall.

Status

This displays the current status and some statistics.
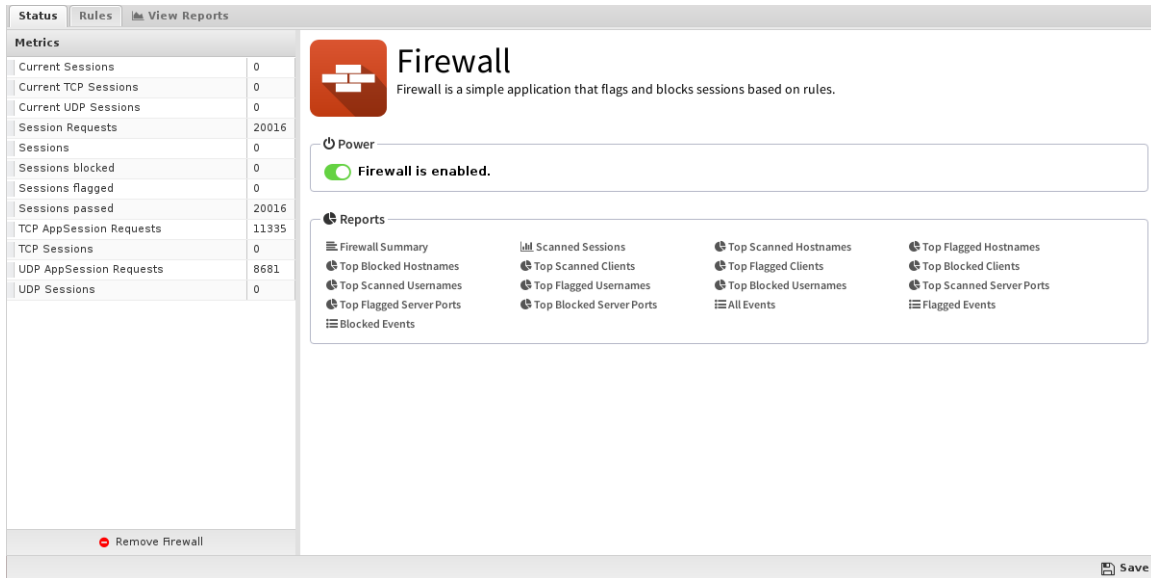
Figure 76: Apps Firewall Status

Rules

The Rules tab allows you to specify rules to Block, Pass or Flag traffic that crosses the NG Firewall.

The Rules documentation describes how rules work and how they are configured. Firewall uses rules to determine to block/pass the specific session, and if the sessions is flagged. Flagging a session marks it in the logs for reviewing in the event logs or reports, but has no direct effect on the network traffic.

Typically NG Firewall is installed as a NAT/gateway device, or behind another NAT/gateway device in bridge mode. In this scenario all inbound sessions are blocked by NAT except those explicitly allowed with port forwards. Because of this, the Firewall does not block anything by default. It is up to you to decide to best fit for your network, whether you only want to block specific ports or you want to block everything and allow only a few services.

Rule Actions

- **Pass**: Allows the traffic which matched the rule to flow.
- **Block**: Blocks the traffic which matched the rule.

Additionally a session can be flagged. If **Flag** is checked the event is flagged for easier viewing in the event log. Flag is always enabled if the action is Block.
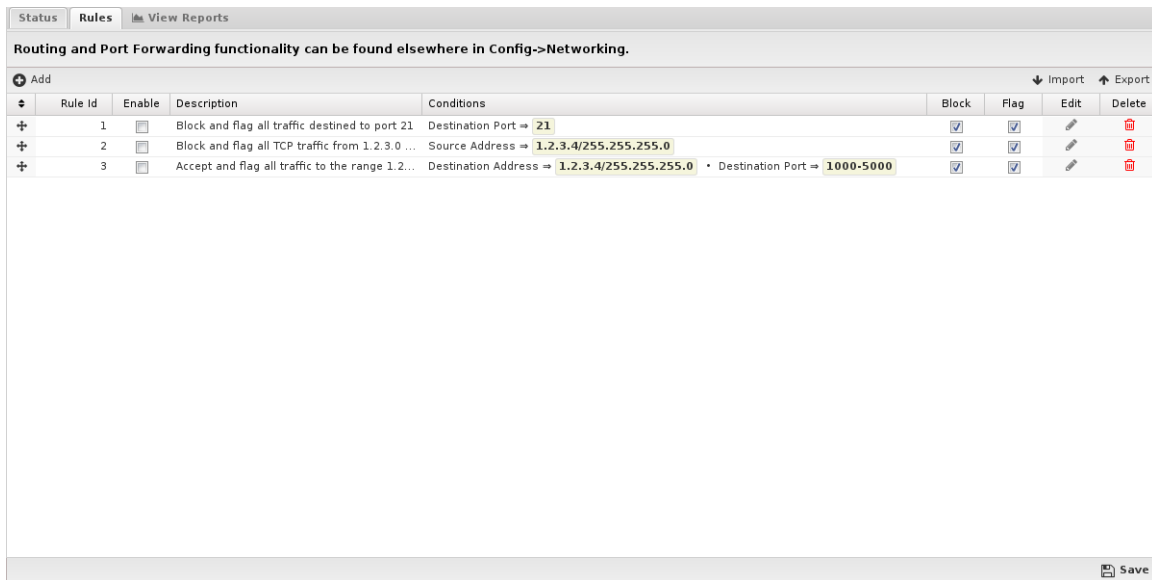
Figure 77: Apps Firewall Rules

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Firewall.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 9: Pre-defined report queries

Report Entry	Description
Firewall Summary	A summary of firewall actions.
Scanned Sessions	The amount of scanned, flagged, and blocked sessions over time.
Top Scanned Hostnames	The number of scanned session grouped by hostname.
Top Flagged Hostnames	The number of flagged session grouped by hostname.
Top Blocked Hostnames	The number of blocked sessions grouped by hostname.
Top Scanned Clients	The number of scanned session grouped by client.
zTop Flagged Clients	The number of flagged session grouped by client.
Top Blocked Clients	The number of blocked session grouped by client.
Top Scanned Usernames	The number of scanned session grouped by username.
Top Flagged Usernames	The number of flagged session grouped by username.
Top Blocked Usernames	The number of blocked session grouped by username.
Top Scanned Server Ports	The number of scanned session grouped by server (destination) port.

Report Entry	Description
Top Flagged Server Ports	The number of flagged session grouped by server (destination) port.
Top Blocked Server Ports	The number of blocked session grouped by server (destination) port.
All Events	All events scanned by Firewall App.
Flagged Events	Events flagged by Firewall App.
Blocked Events	Events blocked by Firewall App.

The tables queried to render these reports:

- Database Schema

Related Topics

NG Firewall User Guide

9.1.1 Firewall Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Firewall.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 10: Pre-defined Report Queries

Report Entry	Description
Firewall Summary	A summary of firewall actions.
Scanned Sessions	The amount of scanned, flagged, and blocked sessions over time.
Top Scanned Hostnames	The number of scanned session grouped by hostname.
Top Flagged Hostnames	The number of flagged session grouped by hostname.
Top Blocked Hostnames	The number of blocked sessions grouped by hostname.
Top Scanned Clients	The number of scanned session grouped by client.
Top Flagged Clients	The number of flagged session grouped by client.
Top Blocked Clients	The number of blocked session grouped by client.
Top Scanned Usernames	The number of scanned session grouped by username.
Top Flagged Usernames	The number of flagged session grouped by username.
Top Blocked Usernames	The number of blocked session grouped by username.
Top Scanned Server Ports	The number of scanned session grouped by server (destination) port.

Report Entry	Description
Top Flagged Server Ports	The number of flagged session grouped by server (destination) port.
Top Blocked Server Ports	The number of blocked session grouped by server (destination) port.
All Events	All events scanned by Firewall App.
Flagged Events	Events flagged by Firewall App.
Blocked Events	Events blocked by Firewall App.

The tables queried to render these reports:

• Database Schema

Related Topics

Reports

Report Viewer

9.2   Intrusion Prevention

Intrusion Prevention is an Intrusion Detection system that detects malicious activity on your network.

To detect malicious activity, Intrusion Prevention uses signatures, a method that draws upon a database of known attack patterns. If a session matches a signature, its enabled action directs Intrusion Prevention to Log (records the incident but does not stop the activity) or Block (records the incident and does stop the activity).

There is tremendous diversity between networks and it is possible for a signature to correctly identify malicious activity on one network and incorrectly match legitimate traffic on another. Logging all matching signatures can make it difficult to effectively monitor Intrusion Prevention, and blocking can disrupt legitimate traffic, causing your network to appear to be broken. Therefore, it is perfectly legitimate for there to be many signatures set as disabled or not active in Intrusion Prevention. In fact, it is advised that you use to the Recommended actions as specified by the signature database providers.

The database contains over 40,000 signatures, making it difficult to manage signatures directly. Rules are used to configure groups of signatures on matching various attributes. A condition can match an attribute such as classtype. For all signatures that match, they are configured in Intrusion Prevention according to the rule action. Any signature not matched by a rule is Disabled. A default set of rules based on system memory are enabled by default.

The signature database is automatically updated several times a week. New and updated rules will be configured as determined by rules.

All detected activity for enabled signatures is recorded to the Intrusion Prevention All Events log. You should review this log on a daily basis.

Note: Intrusion Prevention installs but is off by default.

Note: Intrusion Prevention can be memory intensive and requires at least 2GB of RAM. The amount used is a combination of the number of enabled signatures and the amount of traffic that goes through your system.

Settings

When To Scan

Intrusion Prevention can be run before or after other network processing. Which option depends largely on your reasons for using Intrusion Prevention.

When Before other Network Processing is selected (the default), IPS sees all traffic even if it will subsequently be dropped by the firewall. This means IPS will see much malicious activity like port scan, intrusion attempts on the public IP addresses that happen on almost all networks, even though that traffic will ultimately just be dropped. The advantage of this approach is that Intrusion Prevention sees and logs everything providing the most complete picture. The disadvantage is that it usually logs so much that the Intrusion Prevention event log quickly become ignored because its logging thousands of events per day and this is completely normal and expected.

When After other Network Processing is selected IPS will only scan traffic that will actually pass through the firewall. Most networks where NG Firewall is running with a Public IP and doing NAT and only port forwarding select or no traffic at all, this will be extremely different that scanning "prerouting". The advantage of this mode is that IPS will only scan/log on traffic that is actually entering your network and therefore ignores a lot of the standard "noise" from incoming port scans and vulnerability scans that just get dropped at the firewall and logs only on traffic that should potentially concern the administrator. Another advantage is that it fully allows bypass traffic to work as expected. The disadvantage of this mode is that it provides a less complete picture of activity on the public interface it now no longer logs attempts that just get dropped.

Status

The Status tab shows the following information:

- Memory Usage: The amount of system memory the IPS engine is using compared to your installed system memory.
- Metrics: The number of blocked, logged, and scanned sessions.
- Overview: Signatures and Signature Updates.

  - Signatures: Total number of signatures available and the number set for Log, Block, Disabled.
  - Updates: The last time the signature database was updated and the last time a check was performed. Database updates do not occur on each check.
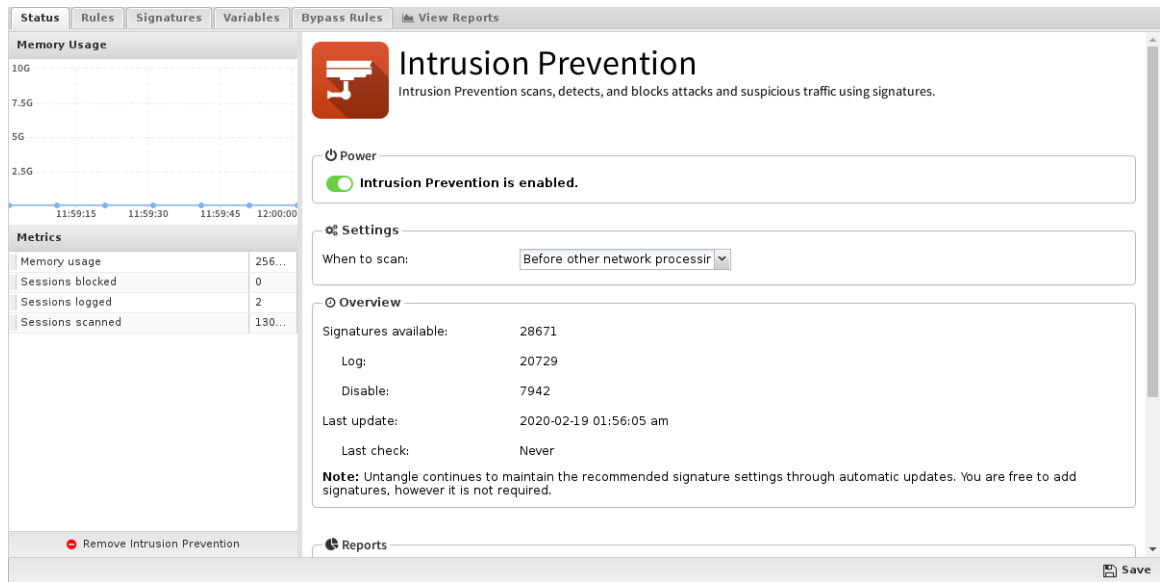
Figure 78: Apps Intrusion-Prevention Status

Rules

Rules allow you to control which signatures are enabled (and their actions) or disabled. For each signature the rules are evaluated in order, the action from the first matching rule is used to determine the status of that signature. The Intrusion Prevention rules are the mechanism to determine which signatures are enabled and their associated actions. These rules have no impact on network traffic and are not evaluated against packets, sessions, or network traffic in any manner.

Any signature not matched by any rule is disabled.

The Rules documentation describes how rules generally work and how they are configured. The major difference for Intrusion Prevention is the Conditions List.

At the bottom of the tab a stats bar indicates how many signatures are affected by the currently defined rules.

When adding or editing a rule, the bottom of the edit window will show how many signatures are affected by the conditions as you build the rule.
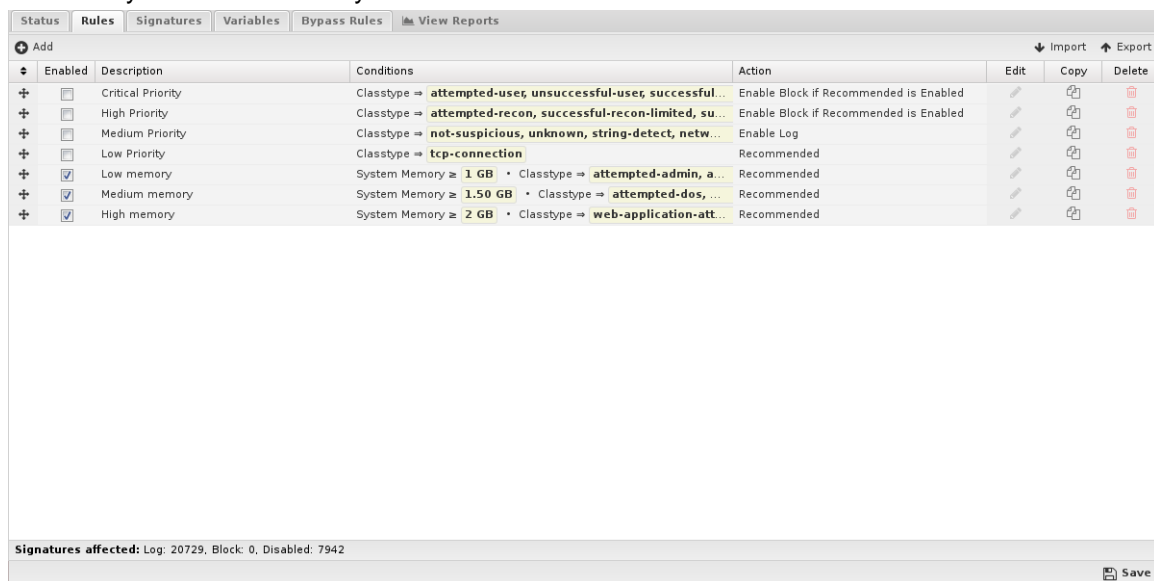
Figure 79: Apps Intrusion-Prevention Rules

Rule Conditions

Conditions define which signatures will match the rule. If and only if all of the conditions match, the rule is considered a match.

The following conditions are specific to Intrusion Prevention rules:

Name	Syntax	Function
Signature identifier	Numeric	Matches if value matches the exact or partial signature identifier.
Group identifier	Numeric	Matches if value matches the exact or partial group identifier.
Category	Checkbox	Matches if value is in one of the checked categories.
Classtype	Checkbox	Matches if value is in one of the checked classtypes.
Message	Text	Matches if value matches the exact or partial signature subject message.
Protocol	Checkbox	Matches if value is in one of the checked protocols.
Source Address	Text	Matches if value matches the exact or partial source address.
Source Port	Text	Matches if value matches the exact or partial source port.
Destination Address	Text	Matches if value matches the exact or partial destination address.
Destination Port	Text	Matches if value matches the exact or partial destination port.
Signature	Text	Matches if value matches the exact or any part of the entire signature.
Custom	Boolean	Matches if value is a custom signature.
Recommended Action	Select	Matches if value is a signature's recommended action.
System Memory	Numeric	Matches if system memory matches this value.

Rule Actions

When all conditions are met, signatures will be configured into Intrusion Prevention as follows:

Action	Function
Recommended	Each signature will use their specific Recommended Action. If that Recommended Action is disabled, it will not be enabled at all.
Enable Log	Each signature will be enabled to log.
Enable Block if Recommended is Enabled	Only if the signature's Recommended Action is Log will the signature be configured for Block. Use this for "wide" condition matches like classtype.

Action	Function
Enable Block	Each signature will be enabled to block. Use this for "narrow" matches like sid and gid.
Disable	Each signature will be disabled and not used by Intrusion Prevention.
Whitelist	Each signature's Source and/or Destination networks will be modified to exclude networks defined by the selected variables.

Signatures

The Signature tab shows the entire database of signatures, both the defaults set provided as well as any custom signatures you may add.

Figure 80: Apps Intrusion-Prevention Signatures

Navigation

By default, signatures are grouped by classtype and you can expand the groups to view the individual signatures.

To better find specific signatures, you can use the Filter to select signature fields and the match you're looking for. The grid view will change to show those signatures matching the filter.

If your filter returned one or more matches, you can create a rule from the filter by clicking **Create Rule**.

Mousing over grid cell will show appropriate information related to that cell. For example, if you mouse over the Rule Action cell, you'll see which rule is affecting this signature.

Custom Signatures

You may create and maintain your own signatures, but most use the default database.

If you wish to add custom signatures you can do so either by clicking **Add**.

Alternatively, if you wish to create a new custom signature on an existing signature, you can click **Copy** then edit that copy.

> **Note:** Don't be tempted to copy a signature to change its Recommended Action. Create a Rule instead!

Variables

This tab provides administrators access to Suricata variables. These variables are used in rules to specify criteria for the source and destination of a packet.

Suricata's most important variable is **$HOME_NET**. **$HOME_NET** defines the network or networks you are trying to protect - it is computer automatically based on your network configuration - it includes all local networks (including aliases). Under nearly every circumstance you will want to leave these values as-is.

Using the **Add** button, custom variables can be added. Adding variables may be used by users adding their own rules.This should only be attempted by advanced users with a strong knowledge of Suricata signature creation.

Figure 81: Apps Intrusion-Prevention Variables

Bypass Rules

Bypass rules enable you to configure traffic that should not be scanned by Intrusion Prevention. The Rules documentation describes how rules generally work and how they are configured.

Figure 82: Apps Intrusion-Prevention Bypass-Rules

Updates

The signature database is checked automatically every night. Updates are typically released 2-3 times week. The signature database does not affect custom signatures. New signatures will be integrated into Intrusion Prevention according to defined rules.

Reports

The Reports tab provides a view of all reports and events for all traffic handled by Intrusion Prevention.

This applications reports can be accessed via the Reports tab at the top or the Reports tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the Conditions window at the bottom of the page. The data used in the report can be obtained on the Current Data window on the right.

Table 11: Pre-defined Report Queries

Report Entry	Description
Intrusion Prevention Summary	A summary of intrusion detection and prevention actions.
Intrusion Detection (all)	The amount of detected and blocked intrusions over time.
Intrusion Detection (logged)	The amount of detected intrusions over time.
Intrusion Detection (blocked)	The amount of blocked intrusions over time.
Top Rules (all)	The number of intrusions detevted by rule.
Top Rules (logged)	The number of intrusions logged by rule.
Top Rules (blocked)	The number of intrusions blocked by rule.
Top Signatures (all)	The number of intrusions detected by signature.
Top Signatures (logged)	The number of intrusions logged by signature.
Top Signatures (blocked)	The number of intrusions blocked by signature.
Top Classtypes (all)	The number of intrusions detected by classtype.
Top Classtypes (logged)	The number of intrusions logged by classtype.
Top Classtypes (blocked)	The number of intrusions blocked by classtype.
Top Categories (all)	The number of intrusions detected by category.
Top Categories (logged)	The number of intrusions logged by category.
Top Categories (blocked)	The number of intrusions blocked by category.
Top Source IP Addresses (all)	The number of intrusions detected by source IP address.
Top Source IP Addresses (logged)	The number of intrusions logged by source IP address.
Top Source IP Addresses (blocked)	The number of intrusions blocked by source IP address.
Top Source Ports (all)	The number of intrusions detected by source port.
Top Source Ports (logged)	The number of intrusions logged by source port.
Top Source Ports (blocked)	The number of intrusions blocked by source port.

Report Entry	Description
Top Destination IP Addresses (all)	The number of intrusions detected by destination IP address.
Top Destination IP Addresses (logged)	The number of intrusions logged by destination IP address.
Top Destination IP Addresses (blocked)	The number of intrusions blocked by destination IP address.
Top Destination Ports (all)	The number of intrusions detected by destination port.
Top Destination Ports (logged)	The number of intrusions logged by destination port.
Top Destination Ports (blocked)	The number of intrusions blocked by destination port.
Top Protocols (all)	The number of intrusions detected by protocol.
Top Protocols (logged)	The number of intrusions logged by protocol.
Top Protocols (blocked)	The number of intrusions blocked by protocol.
All Events	All sessions scanned by Intrusion Prevention.
Logged Events	All sessions matching Intrusion Prevention signatures and logged.
Blocked Events	All sessions matching Intrusion Prevention signatures and blocked.

The tables queried to render these reports:

- Database Schema

All Events

The All Events report shows all enabled signature matches found by Intrusion Prevention.

If there are signatures that are currently set to an action of Log and you determine the signature should in fact be Block, you can click the Block button on the far right. The Block button is disabled for any signature that is already blocked.

Figure 83: Reports Cat Intrusion-prevention Rep All-events

Related Topics

- Intrusion Prevention Systems
- Suricata - Writing Suricata Signatures

9.2.1   Intrusion Prevention Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Intrusion Prevention.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 12: Pre-defined Report Queries

Report Entry	Description
Intrusion Prevention Summary	A summary of intrusion detection and prevention actions.
Intrusion Detection (all)	The amount of detected and blocked intrusions over time.
Intrusion Detection (logged)	The amount of detected intrusions over time.
Intrusion Detection (blocked)	The amount of blocked intrusions over time.

Report Entry	Description
Top Rules (all)	The number of intrusions detevted by rule.
Top Rules (logged)	The number of intrusions logged by rule.
Top Rules (blocked)	The number of intrusions blocked by rule.
Top Signatures (all)	The number of intrusions detected by signature.
Top Signatures (logged)	The number of intrusions logged by signature.
Top Signatures (blocked)	The number of intrusions blocked by signature.
Top Classtypes (all)	The number of intrusions detected by classtype.
Top Classtypes (logged)	The number of intrusions logged by classtype.
Top Classtypes (blocked)	The number of intrusions blocked by classtype.
Top Categories (all)	The number of intrusions detected by category.
Top Categories (logged)	The number of intrusions logged by category.
Top Categories (blocked)	The number of intrusions blocked by category.
Top Source IP Addresses (all)	The number of intrusions detected by source IP address.
Top Source IP Addresses (logged)	The number of intrusions logged by source IP address.
Top Source IP Addresses (blocked)	The number of intrusions blocked by source IP address.
Top Source Ports (all)	The number of intrusions detected by source port.
Top Source Ports (logged)	The number of intrusions logged by source port.
Top Source Ports (blocked)	The number of intrusions blocked by source port.
Top Destination IP Addresses (all)	The number of intrusions detected by destination IP address.
Top Destination IP Addresses (logged)	The number of intrusions logged by destination IP address.
Top Destination IP Addresses (blocked)	The number of intrusions blocked by destination IP address.
Top Destination Ports (all)	The number of intrusions detected by destination port.
Top Destination Ports (logged)	The number of intrusions logged by destination port.
Top Destination Ports (blocked)	The number of intrusions blocked by destination port.
Top Protocols (all)	The number of intrusions detected by protocol.
Top Protocols (logged)	The number of intrusions logged by protocol.
Top Protocols (blocked)	The number of intrusions blocked by protocol.
All Events	All sessions scanned by Intrusion Prevention.
Logged Events	All sessions matching Intrusion Prevention signatures and logged.
Blocked Events	All sessions matching Intrusion Prevention signatures and blocked.

The tables queried to render these reports:

- Database Schema

Related Topics

Report Viewer

Reports

9.3 Phish Blocker

Phish Blocker protects users from phishing attacks over email (SMTP). It inspects email for fraudulent emails, also known as phish. A phishing email attempts to acquire sensitive information such as passwords and credit card details by masquerading as a trustworthy person or business in an apparently official electronic communication, such as an email.

Settings

This section reviews the different settings and configuration options available for Phish Blocker.

Status

This displays the current status and some statistics.

Email

These settings apply only to the scanned **SMTP** messages.

- **Scan SMTP**: This enables or disables **SMTP** scanning.
- **Action**: The action taken on the message if the Spam Score is high enough.

If set to **Mark**, "[Phish]..." will be prepended to the email subject line and it will be delivered. If set to **Pass**, the message will be delivered as originally sent. **Drop** will inform the sending server the mail was successfully delivered, but NG Firewall will drop the mail so it is never delivered. **Quarantine**

will send the mail to users' email quarantine for them to release or delete as they see fit. For more information, refer to Quarantine.

Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Phish Blocker.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 13: Pre-defined Report Queries

Report Entry	Description
Phish Blocker Summary	A summary of phish blocking actions for email activity.
Email Usage (all)	The amount of scanned, clean, and phish email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (clean)	The amount of clean email over time.
Email Usage (phish)	The amount of phish email over time.
Phish Ratio	The ratio of phish (true) to ham (false)
Top Phish Recipients	The number of email addresses with phish.
Top Phish Sender Addresses	The number of IP addresses sending phish.
All Email Events	All email sessions scanned by Phish Blocker.
All Phish Events	All email sessions detected as phishing attempts.
Quarantined Events	All email sessions detected as phishing attempts and quarantined.

The tables queried to render these reports:

- Sessions

- Session Minutes
- Phish Blocker Reports

Related Topics

- Spam Blocker
- Spam Blocker Lite

9.3.1   Phish Blocker Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Phish Blocker.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 14: Pre-defined Report Queries:

Report Entry	Description
Phish Blocker Summary	A summary of phish blocking actions for email activity.
Email Usage (all)	The amount of scanned, clean, and phish email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (clean)	The amount of clean email over time.
Email Usage (phish)	The amount of phish email over time.
Phish Ratio	The ratio of phish (true) to ham (false)
Top Phish Recipients	The number of email addresses with phish.
Top Phish Sender Addresses	The number of IP addresses sending phish.
All Email Events	All email sessions scanned by Phish Blocker.
All Phish Events	All email sessions detected as phishing attempts.
Quarantined Events	All email sessions detected as phishing attempts and quarantined.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Phish Blocker

Related Topics

- Report Viewer
- Reports

9.4   Threat Prevention

Threat Prevention blocks potentially harmful traffic from entering or exiting the network. This app can prevent cyber attacks to your servers (e.g. web, VoIP, and email). It is also useful to prevent data loss in case users mistakenly try to connect to a phishing site or other type of malicious host.

Threat Prevention uses Threat Intelligence technology managed by Webroot BrightCloud®. Webroot BrightCloud® assesses each IP address and provides it a reputation score. The reputation score is the result of running an IP address through BrightCloud's sensor network. The Sensor Network analyzes the IP address based on real time Global Threat Databases that are kept up to date with new and emerging threats. The Threat Prevention app works by performing a query to the BrightCloud® service, requesting for the reputation score and historical data of each IP address or URL. Based on the rating of the IP address or URL, the session may be blocked. By default, the Threat Prevention app blocks sessions with a "High Risk" rating. IP addresses or URLs rated as High Risk may be associated with the following types of attacks:

- **Spam Sources** - IP addresses involved in tunneling spam messages through proxy, anomalous SMTP activities, and forum spam activities.
- **Windows Exploits** - IP addresses participating in the distribution of malware, shell code, rootkits, worms or viruses for Windows platforms.
- **Web Attacks** - IP addresses using cross site scripting, iFrame injection, SQL injection, cross domain injection, or domain password brute force attacks to target vulnerabilities on a web server.
- **Botnets** - IP addresses acting as Botnet Command and Control (C&C) centers, and infected zombie machines controlled by the C&C servers.
- **Denial of Service** - The Denial of Service category includes DOS, DDOS, anomalous sync flood, and anomalous traffic detection.
- **Scanners** - IP addresses involved in unauthorized reconnaissance activities such as probing, host scanning, port scanning and brute force login attempts.
- **Phishing** - IP addresses hosting phishing sites and sites related to other kinds of fraudulent activities.
- **TOR Proxy** - IP addresses acting as exit nodes for the TOR Network. Exit nodes are the last point along the proxy chain and make a direct connection to the originator's intended destination.
- **Proxy** - IP addresses providing proxy services, including both VPN and open web proxy services.
- **Mobile Threats** - Denial of service, packet sniffing, address impersonation, and session hijacking

Settings

This section reviews the different settings and configuration options available for Threat Prevention.

Status

The Status screen shows the running state of Threat Prevention and relevant Metrics such as the number of blocked sessions and high risk threats.
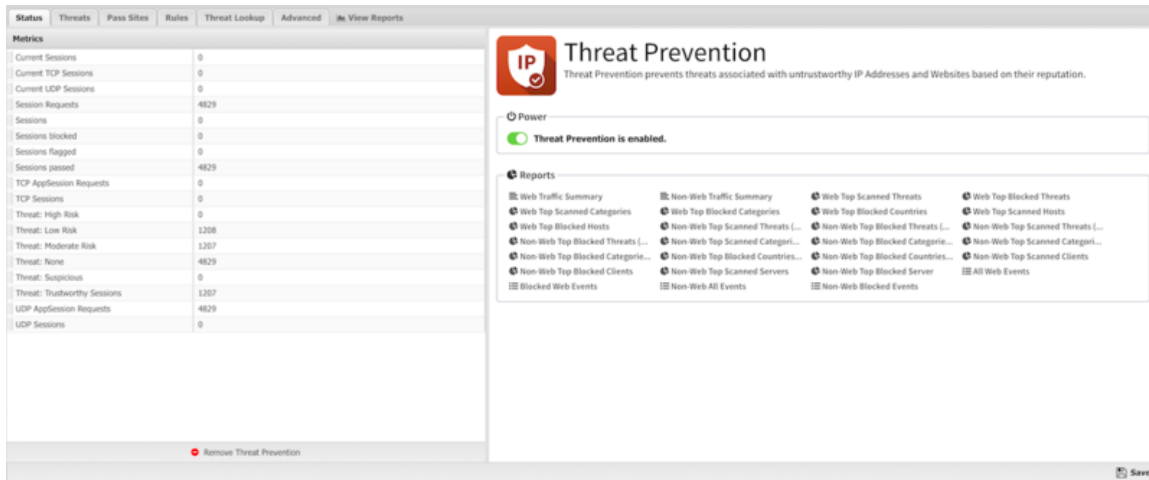
Figure 84: Threat Prevention Status

Threats

In the Threats tab you can review the threshold for IP Addresses and URL Threats. The recommended and default Reputation Threshold is "High Risk". "High Risk" is the only setting that should be deployed without reviewing and understanding the implications on network traffic. 'Suspicious" will block significantly more network traffic than "High Risk" will block.

Figure 85: Threat Prevention Threats

Pass Sites

The **Pass Sites** tab allows you to specify IP Addresses or URLs to exclude from Threat Prevention lookups to ensure they are permitted by this app.
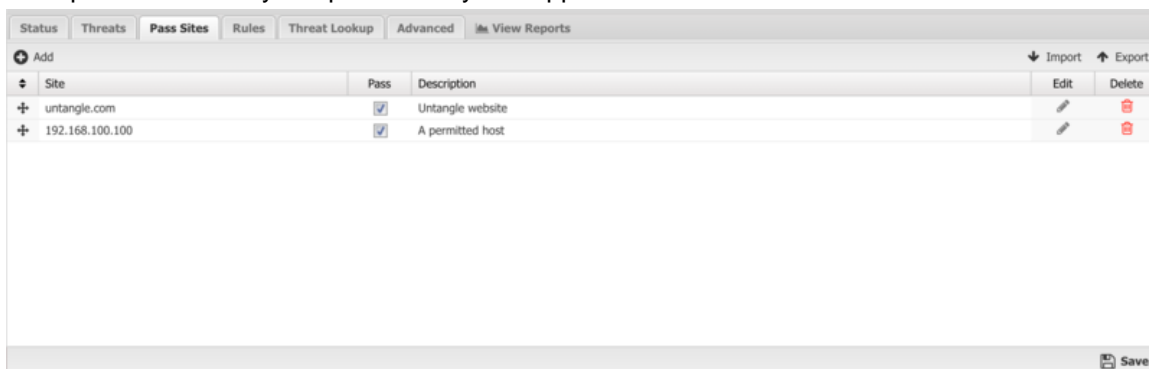
Figure 86: Threat Prevention Pass Sites

Rules

The **Rules** tab allows you to specify rules to Block, Pass or Flag traffic that crosses the NG Firewall.

The Rules describes how rules work and how they are configured. Threat Prevention uses rules to determine to block/pass the specific session, and if the sessions is flagged. Flagging a session marks it in the logs for reviewing in the event logs or reports, but has no direct effect on the network traffic.

In addition to all the common rule types, there are four that are unique to Threat Prevention, and these can be useful for making exceptions to the general *Reputation Threshold* setting.

Client address reputation: The reputation value of a source IP address returned by the Webroot BrightCloud® service. This applies to incoming connections from the Internet to open services on your network.

Server address reputation: The reputation value of a destination IP address returned by the Webroot BrightCloud® service. This applies to outgoing connections to the Internet from hosts on your network.

Client address category: The reputation category of a source IP address returned by the Webroot BrightCloud® service. This applies to incoming connections from the Internet to open services on your network.

Server address category: The reputation category of a destination IP address returned by the Webroot BrightCloud® service. This applies to outgoing connections to the Internet from hosts on your network.

Rule Actions

- **Pass**: Allows the traffic which matched the rule to flow.
- **Block**: Blocks the traffic which matched the rule.

Additionally a session can be flagged. If **Flag** is checked the event is flagged for easier viewing in the event log. Flag is always enabled if the action is Block.
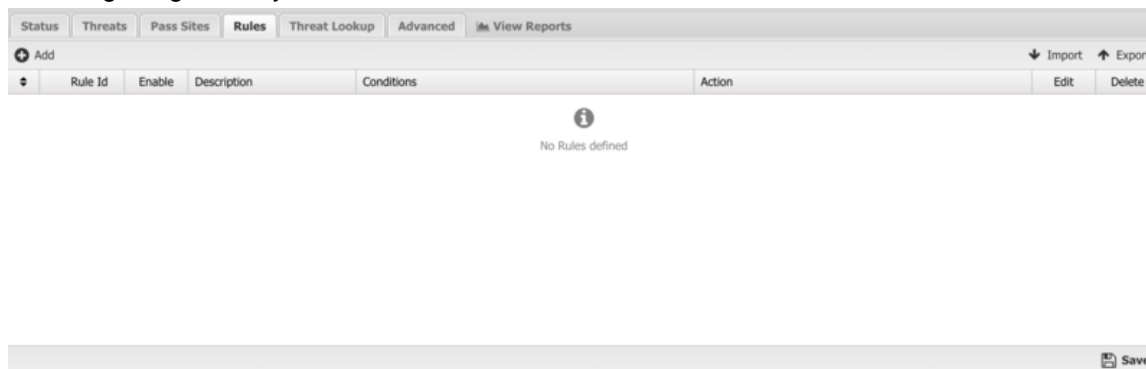
Figure 87: Threat Prevention Rules

Threat Lookup

Threat Lookup enables you to get threat information on an IP Address or URL. This is useful to validate afterwards or confirm in advance the Reputation and other details of the IP Address or URL. Enter an IP Address or URL in the input field and click **Search** to get information.

Figure 88: Threat Prevention Threat Lookup

Threat Results

Result	Description
Address/URL	The IP Address or URL you requested to search.
Country	The country where the IP Address or URL originates.
Popularity	The popularity of the IP Address or URL based on the volume of lookups.
Recent Threat Count	The amount of recent occurrences that the IP Address or URL has been associated to a threat.
Age	The amount of time since the IP Address or URL was first noticed.
Reputation	The reputation of the IP Address or URL as determined by the Webroot BrightCloud reputation service.
Details	A description of the Reputation value.

Advanced

The Advanced section enables you to configure additional Threat Prevention options.

Custom block page URL: Set an external location where to redirect users when they are denied access to a web site by Threat Prevention. This is useful if you want to use your own server to process the denial in a different way than the built-in denial options.

Enabling this option will only redirect internal/outbound traffic to your custom page. It will not function to redirect external/inbound traffic (such as Port Forwarded traffic).

Block Options: *Close connection for blocked HTTPS sessions without redirecting to block page* If enabled, secure sites blocked by Threat Prevention do not redirect the user to a denial page and close the connection without any notice to the user. This is useful when you are not using SSL Inspector and the server's root certificate is not installed on the client device.

Figure 89: Threat Prevention Advanced

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions***Conditions* window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Web Traffic Summary	A summary of web Threat Prevention actions.
Non-Web Traffic Summary	A summary of non-web Threat Prevention actions.
Web Top Scanned Threats	The number of web scanned sessions to servers grouped by threat reputation.
Web Top Blocked Threats	The number of web blocked sessions to servers grouped by threats reputation.
Web Top Scanned Categories	The number of other scanned sessions to servers grouped by threat.
Web Top Blocked Categories	The number of web sessions blocked grouped by threat.
Web Top Blocked Countries	Top blocked web sessions to servers grouped by country.
Web Top Scanned Hosts	The number of web scanned sessions grouped by server.
Web Top Blocked Hosts	The number of web blocked session grouped by client.
Non-Web Top Scanned Threats (by client)	The number of non-web scanned sessions from clients grouped by threat reputation.
Non-Web Top Blocked Threats (by client)	The number of non-web blocked sessions from clients grouped by threat reputation.
Non-Web Top Scanned Threats (by server)	The number of non-web scanned sessions to servers grouped by threat reputation.
Non-Web Top Blocked Threats (by server)	The number of non-web blocked sessions to servers grouped by threat reputation.
Non-Web Top Scanned Categories (by client)	The number of non-web scanned sessions from clients grouped by threat.

Report Entry	Description
Non-Web Top Blocked Categories (by client)	The number of non-web blocked sessions from clients grouped by threat.
Non-Web Top Scanned Categories (by server)	The number of non-web scanned sessions to servers grouped by threat.
Non-Web Top Blocked Categories (by server)	The number of non-web blocked sessions to servers grouped by threat.
Non-Web Top Blocked Countries (by client)	Top non-web blocked sessions from clients grouped by country.
Non-Web Top Blocked Countries (by server)	Top non-web blocked sessions to servers grouped by threat.
Non-Web Top Scanned Clients	The number of non-web scanned session grouped by client.
Non-Web Top Blocked Clients	The number of non-web blocked session grouped by client.
Non-Web Top Scanned Servers	The number of non-web scanned sessions grouped by server.
Non-Web Top Blocked Server	The number of non-web blocked session grouped by client.
All Web Events	Shows all scanned web requests.
Blocked Web Events	Shows all blocked web requests.
Non-Web All Events	All non-web events scanned by Threat Prevention.
Non-Web Blocked Events	Non-web events blocked by Threat Prevention.

The tables queried to render these reports:

- Sessions
- Session Minutes

9.4.1   Threat Prevention Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports

Reports can be searched and further defined using the time selectors and the **Condition** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Web Traffic Summary	A summary of web Threat Prevention actions.
Non-Web Traffic Summary	A summary of non-web Threat Prevention actions.
Web Top Scanned Threats	The number of web scanned sessions to servers grouped by threat reputation.
Web Top Blocked Threats	The number of web blocked sessions to servers grouped by threats reputation.

Report Entry	Description
Web Top Scanned Categories	The number of other scanned sessions to servers grouped by threat.
Web Top Blocked Categories	The number of web sessions blocked grouped by threat.
Web Top Blocked Countries	Top blocked web sessions to servers grouped by country.
Web Top Scanned Hosts	The number of web scanned sessions grouped by server.
Web Top Blocked Hosts	The number of web blocked session grouped by client.
Non-Web Top Scanned Threats (by client)	The number of non-web scanned sessions from clients grouped by threat reputation.
Non-Web Top Blocked Threats (by client)	The number of non-web blocked sessions from clients grouped by threat reputation.
Non-Web Top Scanned Threats (by server)	The number of non-web scanned sessions to servers grouped by threat reputation.
Non-Web Top Blocked Threats (by server)	The number of non-web blocked sessions to servers grouped by threat reputation.
Non-Web Top Scanned Categories (by client)	The number of non-web scanned sessions from clients grouped by threat.
Non-Web Top Blocked Categories (by client)	The number of non-web blocked sessions from clients grouped by threat.
Non-Web Top Scanned Categories (by server)	The number of non-web scanned sessions to servers grouped by threat.
Non-Web Top Blocked Categories (by server)	The number of non-web blocked sessions to servers grouped by threat.
Non-Web Top Blocked Countries (by client)	Top non-web blocked sessions from clients grouped by country.
Non-Web Top Blocked Countries (by server)	Top non-web blocked sessions to servers grouped by threat.
Non-Web Top Scanned Clients	The number of non-web scanned session grouped by client.
Non-Web Top Blocked Clients	The number of non-web blocked session grouped by client.
Non-Web Top Scanned Servers	The number of non-web scanned sessions grouped by server.
Non-Web Top Blocked Server	The number of non-web blocked session grouped by client.
All Web Events	Shows all scanned web requests.
Blocked Web Events	Shows all blocked web requests.
Non-Web All Events	All non-web events scanned by Threat Prevention.
Non-Web Blocked Events	Non-web events blocked by Threat Prevention.

The tables queried to render these reports:

- Sessions
- Session Minutes

9.5   Virus Blocker

Virus Blocker transparently scans your HTTP, FTP, and SMTP traffic to protect your network from viruses, trojans and other malware. It scans within archives such as zip, rar, tar, gzip, bzip2 (and more).

As files are downloaded onto the network Virus Blocker scans downloads using many technologies. Firstly, it will collect metadata about the file and query the NG Firewall threat intelligence database for information about the file based on its fingerprint. Secondly, a local scan using Bitdefender's signature database will be run on the server while the cloud lookup is being performed. Thirdly, a heuristic scan is performed to look for suspicious patterns in executable files. Lastly, dynamic analysis is performed by evaluating code in a emulator and looking for malicious activity. If the download fails any of the above test it is considered malware and the download is blocked.

Settings

This section reviews the different settings and configuration options available for the virus scanners.

Status

This displays the current status and some statistics.

Figure 90: Virus Blocker Status

Scan Options

Scan options configure what network traffic and content to scan.

- **Scan HTTP**: This enables or disables HTTP scanning.
- **Scan SMTP**: This option enables scanning of SMTP **message attachments**.
- **Action**: The selected action will be taken on a message if a virus is found.

Setting Action to **Remove Infection** will remove the infected attachment and wrap the original email for delivery to the intended recipient. If set to **Pass Message**, the original message will be wrapped and delivered with the attachment intact. In both cases, the subject line is prepended with "[VIRUS]". **Block** will block the message from being delivered.

- **Scan FTP**: This enables or disables scanning of FTP downloads.

Figure 91: Virus Blocker Scan Options

Pass Sites

This section allows you to specify sites that are not scanned. The list uses the Glob Matcher syntax.

Note: Use caution when adding sites to this list!

For each protocol, the behavior is as follows:

- **HTTP**: Match the HTTP Host header.
- **FTP**: Match the server IP address or domain address (if a reverse DNS address exists).
- **Email**: Match the client or server IP address or domain address (if a reverse DNS address exists).
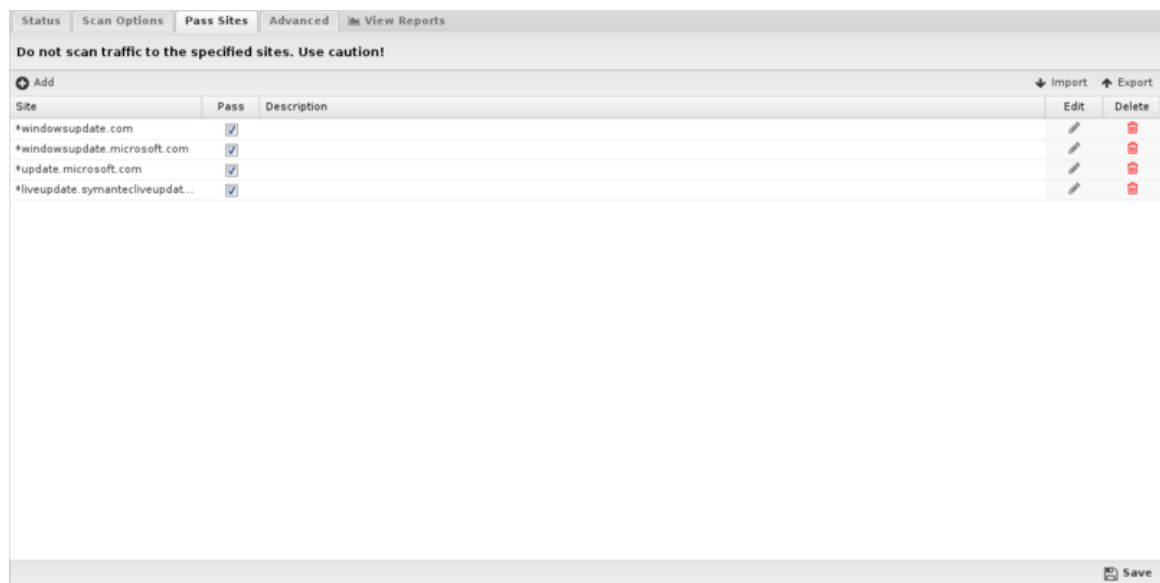
Figure 92: Virus Blocker Pass Sites

Advanced

Advanced settings can tune specific behavior of virus blocker.

The first options can enable/disable certain scanners. When a file is scanned by virus blocker, it is scanned by multiple engines, a local antivirus engine and the cloud ScoutIQ™ engine.

Using all available engines is recommended.
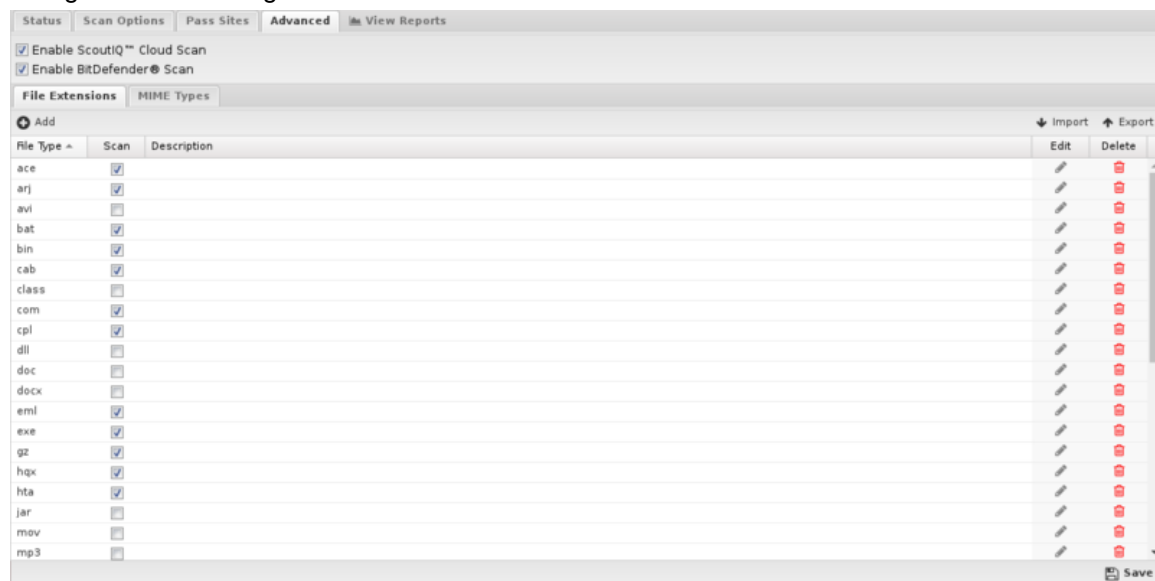
Figure 93: Virus Blocker Advanced

File Extensions

File extensions configure which HTTP files will be scanned. The defaults are the recommended values. However, in some cases you may with to add or remove certain file extensions.

A understanding of security tradeoffs and pragmatism is essential before changing these settings. Unlike the other URL-based scanning of other apps like Web Filter, Virus Blocker runs in depth analysis of the file, including signatures, heuristics, and emulation. Unlike host-based antivirus, the

gateway is a unique resource shared among the whole network and furthermore it has no ability to scan-on-exec as it has no knowledge of what the client plans to execute, unlike host-based antivirus. Scanning is expensive and turning on certain extensions (like .png files) can cripple the network. Analyzing reports to see how many scans are being done and if those resources are being spent on worthwhile scan resources is a good exercise. It is not uncommon to see millions of scans of some application's update.

MIME Types

Similar to file extensions, but this lists the MIME types to be scanned, regardless of extension. The same logic and warnings apply here as well.

Reports

The **Reports**tab provides a view of all reports and events for all traffic handled by Virus Blocker.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Virus Blocker Web Summary	A summary of virus blocking actions for web activity.
Virus Blocker FTP Summary	A summary of virus blocking actions for FTP activity.
Virus Blocker Email Summary	A summary of virus blocking actions for Email activity.
Web Usage (all)	The amount of scanned and blocked web requests over time.
Web Usage (scanned)	The amount of scanned web requests over time.
Web Usage (blocked)	The amount of blocked web requests over time.
Web Top Blocked Viruses	The top web virus blocked.
Web Top Blocked Clients	The top web clients by blocked virus count.
Web Top Blocked Sites	The top web sites by blocked virus count.
Web Top Scanned Sites	The top web sites by scan count.
FTP Usage (all)	The amount of scanned and blocked FTP requests over time.
FTP Usage (scanned)	The amount of scanned FTP requests over time.
FTP Usage (blocked)	The amount of blocked FTP requests over time.
FTP Top Blocked Viruses	The number of blocked viruses by FTP activity.
FTP Top Blocked Clients	The number of clients with blocked viruses by FTP activity.
FTP Top Blocked Sites	The number of clients with blocked viruses by FTP activity.
Email Usage (all)	The amount of scanned and blocked email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (blocked)	The amount of blocked email over time.
Email Top Blocked Viruses	The number of blocked viruses by Email activity.

Report Entry	Description
Email Top Blocked Clients	The number of clients with blocked viruses by Email activity.
Email Top Blocked Sites	The number of clients with blocked viruses by Email activity.
Scanned Web Events	All HTTP sessions scanned by Virus Blocker.
Infected Web Events	Infected HTTP sessions blocked by Virus Blocker.
Clean Web Events	Scanned HTTP sessions marked clean.
Scanned Email Events	All email sessions scanned by Virus Blocker.
Infected Email Events	Infected email sessions blocked by Virus Blocker.
Clean Email Events	Scanned email sessions marked clean.
Scanned FTP Events	All FTP sessions scanned by Virus Blocker.
Infected FTP Events	Infected FTP sessions blocked by Virus Blocker.
Clean FTP Events	Scanned FTP sessions marked clean.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Virus Blocker Reports

Related Topics

Virus Blocker Lite

9.5.1 Virus Blocker Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Virus Blocker.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Pre-defined report queries:

Report Entry	Description
Virus Blocker Web Summary	A summary of virus blocking actions for web activity.
Virus Blocker FTP Summary	A summary of virus blocking actions for FTP activity.
Virus Blocker Email Summary	A summary of virus blocking actions for Email activity.
Web Usage (all)	The amount of scanned and blocked web requests over time.
Web Usage (scanned)	The amount of scanned web requests over time.
Web Usage (blocked)	The amount of blocked web requests over time.
Web Top Blocked Viruses	The top web virus blocked.
Web Top Blocked Clients	The top web clients by blocked virus count.

Report Entry	Description
Web Top Blocked Sites	The top web sites by blocked virus count.
Web Top Scanned Sites	The top web sites by scan count.
FTP Usage (all)	The amount of scanned and blocked **FTP** requests over time.
FTP Usage (scanned)	The amount of scanned FTP requests over time.
FTP Usage (blocked)	The amount of blocked FTP requests over time.
FTP Top Blocked Viruses	The number of blocked viruses by FTP activity.
FTP Top Blocked Clients	The number of clients with blocked viruses by FTP activity.
FTP Top Blocked Sites	The number of clients with blocked viruses by FTP activity.
Email Usage (all)	The amount of scanned and blocked email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (blocked)	The amount of blocked email over time.
Email Top Blocked Viruses	The number of blocked viruses by Email activity.
Email Top Blocked Clients	The number of clients with blocked viruses by Email activity.
Email Top Blocked Sites	The number of clients with blocked viruses by Email activity.
Scanned Web Events	All HTTP sessions scanned by Virus Blocker.
Infected Web Events	Infected HTTP sessions blocked by Virus Blocker.
Clean Web Events	Scanned HTTP sessions marked clean.
Scanned Email Events	All email sessions scanned by Virus Blocker.
Infected Email Events	Infected email sessions blocked by Virus Blocker.
Clean Email Events	Scanned email sessions marked clean.
Scanned FTP Events	All FTP sessions scanned by Virus Blocker.
Infected FTP Events	Infected FTP sessions blocked by Virus Blocker.
Clean FTP Events	Scanned FTP sessions marked clean.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Virus Blocker

Related Topics

Report Viewer

Reports

9.6   Virus Blocker Lite

Virus Blocker Lite transparently scans your HTTP, FTP, and SMTP traffic to protect your network from viruses, trojans and other malware. It scans within archives such as zip, rar, tar, gzip, bzip2 (and more).

Virus Blocker Lite is based on an open-source virus scanner, Clam AV. Clam AV is well-known for its speed and accuracy.

Settings

This section reviews the different settings and configuration options available for the virus scanners.

Status

This displays the current status and some statistics.

Scan Options

Scan options configure what network traffic and content to scan.

- **Scan HTTP**: This enables or disables HTTP scanning.
- **Scan SMTP**: This option enables scanning of SMTP **message attachments**.
- **Action**: The selected action will be taken on a message if a virus is found.

  - Setting Action to **Remove Infection** will remove the infected attachment and wrap the original email for delivery to the intended recipient.
  - If set to **Pass Message**, the original message will be wrapped and delivered with the attachment intact.

    Note: In both cases, the subject line is prepended with **[VIRUS]**.

- **Block** will block the message from being delivered.
- **Scan FTP**: This enables or disables scanning of FTP downloads.

Pass Sites

This section allows you to specify sites that are not scanned. The list uses the Glob Matcher syntax.

Note: Use caution when adding sites to this list!

For each protocol, the behavior is as follows:

- **HTTP**: Match the HTTP Host header.
- **FTP**: Match the server IP address or domain address (if a reverse DNS address exists).
- **Email**: Match the client or server IP address or domain address (if a reverse DNS address exists).

Advanced

Advanced settings can tune specific behavior of virus blocker.

The first options can enable/disable certain scanners. When a file is scanned by virus blocker, it is scanned by multiple engines, a local antivirus engine and the cloud ScoutIQ™ engine.

Using all available engines is recommended.

File Extensions

File extensions configure which HTTP files will be scanned. The defaults are the recommended values. However, in some cases you may with to add or remove certain file extensions.

A understanding of security tradeoffs and pragmatism is essential before changing these settings. Unlike the other URL-based scanning of other apps like Web Filter, Virus Blocker runs in depth analysis of the file, including signatures, heuristics, and emulation. Unlike host-based antivirus, the gateway is a unique resource shared among the whole network and furthermore it has no ability to scan-on-exec as it has no knowledge of what the client plans to execute, unlike host-based antivirus. Scanning is expensive and turning on certain extensions (like .png files) can cripple the network. Analyzing reports to see how many scans are being done and if those resources are being spent on worthwhile scan resources is a good exercise. It is not uncommon to see millions of scans of some application's update.

MIME Types

Similar to file extensions, but this lists the MIME types to be scanned, regardless of extension. The same logic and warnings apply here as well.

Reports

The Reports tab provides a view of all reports and events for all traffic handled by Virus Blocker Lite.

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 15: Pre-defined Report Queries

Report Entry	Description
Virus Blocker Lite Web Summary	A summary of virus blocking actions for web activity.
Virus Blocker Lite FTP Summary	A summary of virus blocking actions for FTP activity.
Virus Blocker Lite Email Summary	A summary of virus blocking actions for Email activity.

Report Entry	Description
Web Usage (all)	The amount of scanned and blocked web requests over time.
Web Usage (scanned)	The amount of scanned web requests over time.
Web Usage (blocked)	The amount of blocked web requests over time.
Web Top Blocked Viruses	The top web virus blocked.
Web Top Blocked Clients	The top web clients by blocked virus count.
Web Top Blocked Sites	The top web sites by blocked virus count.
Web Top Scanned Sites	The top web sites by scan count.
FTP Usage (all)	The amount of scanned and blocked FTP requests over time.
FTP Usage (scanned)	The amount of scanned FTP requests over time.
FTP Usage (blocked)	The amount of blocked FTP requests over time.
FTP Top Blocked Viruses	The number of blocked viruses by FTP activity.
FTP Top Blocked Clients	The number of clients with blocked viruses by FTP activity.
FTP Top Blocked Sites	The number of clients with blocked viruses by FTP activity.
Email Usage (all)	The amount of scanned and blocked email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (blocked)	The amount of blocked email over time.
Email Top Blocked Viruses	The number of blocked viruses by Email activity.
Email Top Blocked Clients	The number of clients with blocked viruses by Email activity.
Email Top Blocked Sites	The number of clients with blocked viruses by Email activity.
Scanned Web Events	All HTTP sessions scanned by Virus Blocker Lite.
Infected Web Events	Infected HTTP sessions blocked by Virus Blocker Lite.
Clean Web Events	Scanned HTTP sessions marked clean.
Scanned Email Events	All email sessions scanned by Virus Blocker Lite.
Infected Email Events	Infected email sessions blocked by Virus Blocker Lite.
Clean Email Events	Scanned email sessions marked clean.
Scanned FTP Events	All FTP sessions scanned by Virus Blocker Lite.
Infected FTP Events	Infected FTP sessions blocked by Virus Blocker Lite.
Clean FTP Events	Scanned FTP sessions marked clean.

The tables queried to render these reports:

- Sessions
- Session Minutes

- Virus Blocker Lite Reports

Related Topics

Virus Blocker

9.6.1   Virus Blocker Lite Reports

The **Reports** tab provides a view of all reports and events for all traffic handled by Virus Blocker Lite.

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 16: Pre-Defined Report Queries

Report Entry	Description
Virus Blocker Lite Web Summary	A summary of virus blocking actions for web activity.
Virus Blocker Lite FTP Summary	A summary of virus blocking actions for FTP activity.
Virus Blocker Lite Email Summary	A summary of virus blocking actions for Email activity.
Web Usage (all)	The amount of scanned and blocked web requests over time.
Web Usage (scanned)	The amount of scanned web requests over time.
Web Usage (blocked)	The amount of blocked web requests over time.
Web Top Blocked Viruses	The top web virus blocked.
Web Top Blocked Clients	The top web clients by blocked virus count.
Web Top Blocked Sites	The top web sites by blocked virus count.
Web Top Scanned Sites	The top web sites by scan count.
FTP Usage (all)	The amount of scanned and blocked FTP requests over time.
FTP Usage (scanned)	The amount of scanned FTP requests over time.
FTP Usage (blocked)	The amount of blocked FTP requests over time.
FTP Top Blocked Viruses	The number of clients with blocked viruses by FTP activity.
FTP Top Blocked Clients	The number of clients with blocked viruses by FTP activity.
FTP Top Blocked Sites	The number of clients with blocked viruses by FTP activity.
Email Usage (all)	The amount of scanned and blocked email over time.
Email Usage (scanned)	The amount of scanned email over time.
Email Usage (blocked)	The amount of blocked email over time.

Report Entry	Description
Email Top Blocked Viruses	The number of blocked viruses by Email activity.
Email Top Blocked Clients	The number of clients with blocked viruses by Email activity.
Email Top Blocked Sites	The number of clients with blocked viruses by Email activity.
Scanned Web Events	All HTTP sessions scanned by Virus Blocker Lite.
Infected Web Events	Infected HTTP sessions blocked by Virus Blocker Lite.
Clean Web Events	Scanned HTTP sessions marked clean.
Scanned Email Events	All email sessions scanned by Virus Blocker Lite.
Infected Email Events	Infected email sessions blocked by Virus Blocker Lite.
Clean Email Events	Scanned email sessions marked clean.
Scanned FTP Events	All FTP sessions scanned by Virus Blocker Lite.
Infected FTP Events	Infected FTP sessions blocked by Virus Blocker Lite.
Clean FTP Events	Scanned FTP sessions marked clean.

The tables queried to render these reports:

- Sessions
- Session Minutes
- Virus Blocker Lite

Related Topics

Report Viewer

Reports

9.7   Virus Blockers Common

This section reviews the different settings and configuration options available for the virus scanners.

Status

This displays the current status and some statistics.

Scan Options

Scan options configure what network traffic and content to scan.

- **Scan HTTP**: This enables or disables HTTP scanning.
- **Scan SMTP**: This option enables scanning of SMTP **message attachments**.
- **Action**: The selected action will be taken on a message if a virus is found.

  Setting Action to **Remove Infection** will remove the infected attachment and wrap the original email for delivery to the intended recipient. If set to **Pass Message**, the original message will be wrapped and delivered with the attachment intact. In both cases, the subject line is prepended with "[VIRUS]". **Block** will block the message from being delivered.
- **Scan FTP**: This enables or disables scanning of FTP downloads.

Pass Sites

This section allows you to specify sites that are not scanned. The list uses the Glob Matcher syntax.

> ✎ **Note:** Use caution when adding sites to this list!

For each protocol, the behavior is as follows:

- **HTTP**: Match the HTTP Host header.
- **FTP**: Match the server IP address or domain address (if a reverse DNS address exists).
- **Email**: Match the client or server IP address or domain address (if a reverse DNS address exists).

Advanced

Advanced settings can tune specific behavior of virus blocker.

The first options can enable/disable certain scanners. When a file is scanned by virus blocker, it is scanned by multiple engines, a local antivirus engine and the cloud ScoutIQ™ engine.

Using all available engines is recommended.

File Extensions

File extensions configure which HTTP files will be scanned. The defaults are the recommended values. However, in some cases you may with to add or remove certain file extensions.

A understanding of security tradeoffs and pragmatism is essential before changing these settings. Unlike the other URL-based scanning of other apps like Web Filter, Virus Blocker runs in depth

analysis of the file, including signatures, heuristics, and emulation. Unlike host-based antivirus, the gateway is a unique resource shared among the whole network and furthermore it has no ability to scan-on-exec as it has no knowledge of what the client plans to execute, unlike host-based antivirus. Scanning is expensive and turning on certain extensions (like .png files) can cripple the network. Analyzing reports to see how many scans are being done and if those resources are being spent on worthwhile scan resources is a good exercise. It is not uncommon to see millions of scans of some application's update.

MIME Types

Similar to file extensions, but this lists the MIME types to be scanned, regardless of extension. The same logic and warnings apply here as well.

NG Firewall Additional Apps

Contents

10.1   Configuration Backup

NG Firewall's Configuration Backup enables you to recover from hardware failures and disasters. Configuration backup is also used to replicate configuration across multiple deployments of NG Firewall. If installed and enabled, Configuration Backup automatically backs up your configuration on a daily basis to Dashboard and as a secondary option to Google Drive.

Prerequisites

- To use Configuration Backup you must have the Complete Package or Live Reports , which is available as an à la carte item.
- To view the status of backups you need the Reports app.
- To access your backups your appliance must be connected to your ETM Dashboard account. To verify that NG Firewall is connected to ETM Dashboard refer to the Support screen.
- To backup to Google Drive, a Google account must be authenticated using Directory Connector .

Installing the Configuration Backup App

To enable automatic backup you must first install the Configuration Backup app.

1. In the NG Firewall administration, click **Apps** in the menu at the top of the screen.
2. Verify if the **Configuration Backup** app is installed. If not, follow the steps below to install the app.
3. Click **Install Apps**.
4. Click the **Configuration Backup** app.
5. Click **Back to Apps** and wait for the app to finish the installation.

Configuring automatic backup

1. In the **Apps** screen, click the **Configuration Backup** app to configure backups.
2. Toggle the **Power** switch to enable or disable automatic backups.
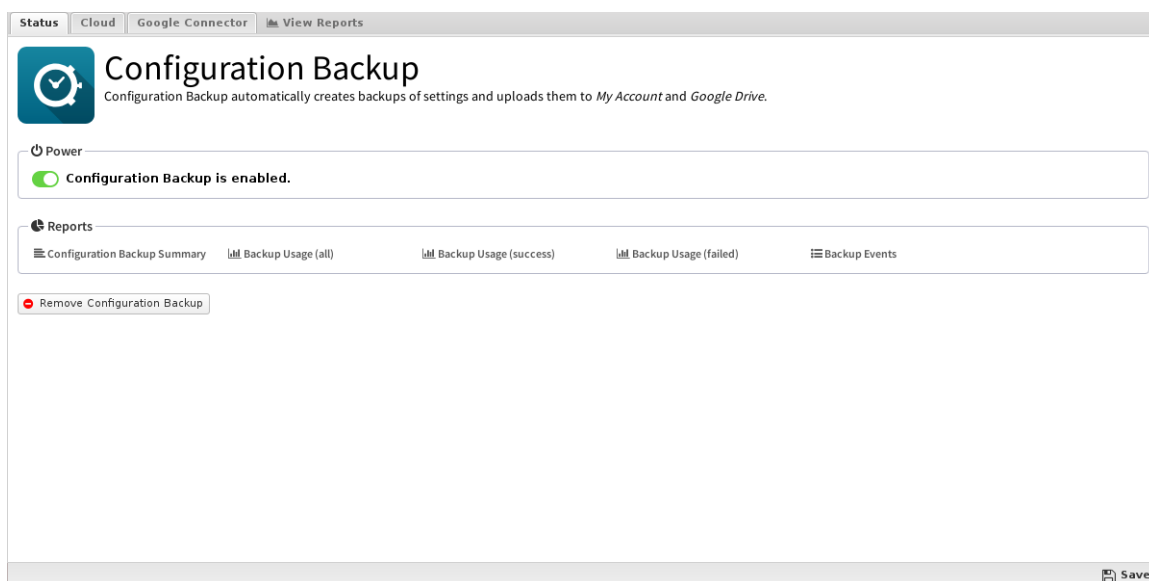
3. Click the **Cloud** tab.
4. Review the Daily Backup setting and modify the Hour and Minute if necessary based on when you want the daily backup to occur.
5. Click **Backup now** if you want to initiate an immediate backup.

Configuring a Secondary Backup using Google Drive

You can use Google Drive as a secondary backup option. Before configuring backups to Google drive, you must connect Google your account in NG Firewall. To configure backup to Google Drive:

1. Go to the **Google Connector** tab.
2. Check **Enable upload to Google Drive:**
3. Confirm the name of your **Google Drive Directory**.
4. Click **Save**.
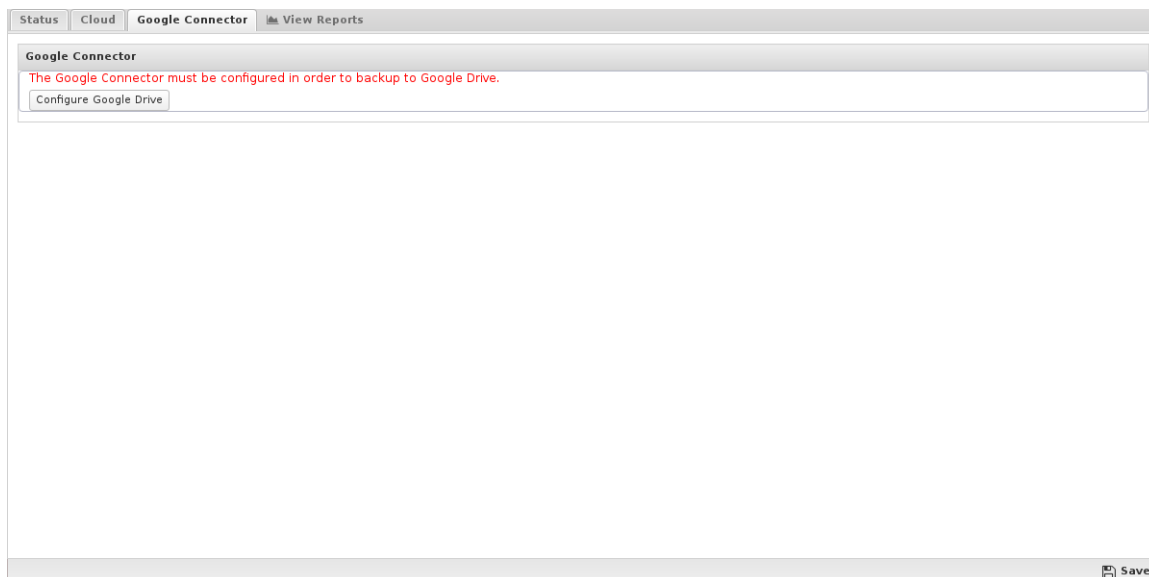
Viewing Backup Activity

To view backup activities such as the most recent backup or potential failures, click **Reports** in the top menu, or click one of the predefined reports in the **Status** screen.

Reports

The **Configuration Backup** reports summarize the backup activities including successful or failed backups.

Configuration Backup / **Configuration Backup Summary**
A summary of configuration backup actions.

Configuration backup attempted 0 backups where 0 succeeded and 0 failed.

Table 17: Pre-Defined Report Queries

Report Entry	Description
Configuration Backup Summary	A summary of configuration backup actions.
Backup Usage (all)	The amount of successes, and failures of configuration backup over time.
Backup Usage (success)	The amount of successful configuration backups over time.
Backup Usage (failed)	The amount of failed configuration backups over time.
Backup Events	All Configuration Backup events.

Restoring a backup using ETM Dashboard

To restore a backup from ETM Dashboard:

1. Log in to Dashboard with your account.
2. Click **Appliances** in the top menu.
3. Select an appliance from the **Appliances** list.
4. Locate the **Cloud Backups** panel and select a backup based on the timestamp.
5. Click **Download**.
6. Click **Yes** to confirm.

Restoring a Backup in Google Drive
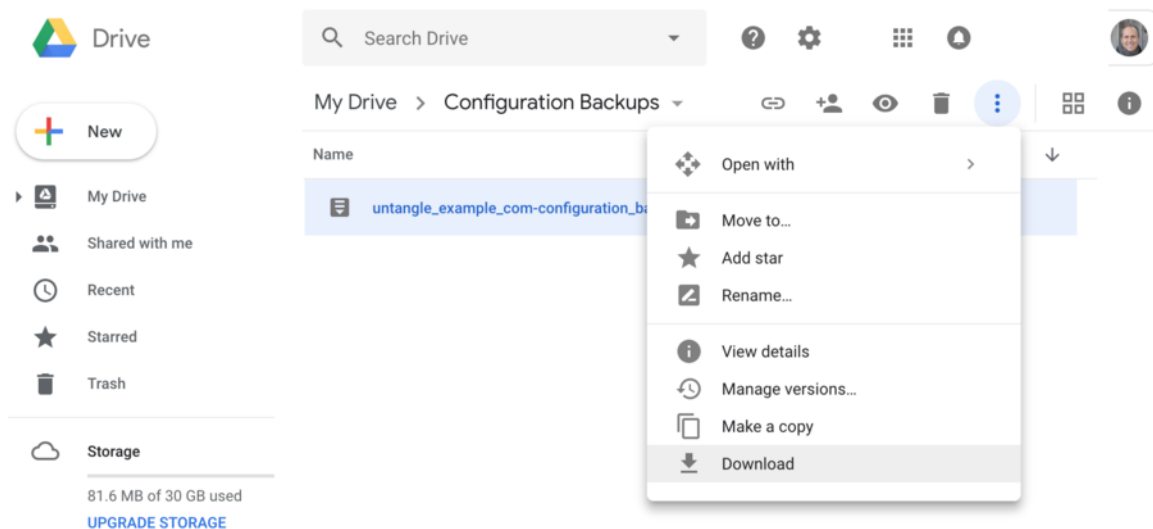
To restore a backup from Google Drive:

1. Log in to your Google account.
2. Go to your Google Drive and locate the directory used by NG Firewall (e.g. Configuration Backups).
3. Select the backup file and click **Download** from the menu.
4. Log in to your NG Firewall.
5. Go to **config** > **System** > **Restore**.
6. Choose a **Restore Option**.
7. Click **Restore from File** and select your backup file.

Related Topics

Restore

10.1.1   Configuration Backup Reports

The **Reports** tab provides a view of all reports and events for Configuration Backup

Reports

This applications reports can be accessed via the **Reports** tab at the top or the **Reports** tab within the settings. All pre-defined reports will be listed along with any custom reports that have been created.

Reports can be searched and further defined using the time selectors and the **Conditions** window at the bottom of the page. The data used in the report can be obtained on the **Current Data** window on the right.

Table 18: Pre-Defined Report Queries

Report Entry	Description
Configuration Backup Summary	A summary of configuration backup actions.
Backup Usage (all)	The amount of successes, and failures of configuration backup over time.
Backup Usage (success)	The amount of successful configuration backups over time.
Backup Usage (failed)	The amount of failed configuration backups over time.
Backup Events	All Configuration Backup events.

Related Topics

Report Viewer

Reports

10.2   Live Support

Live Support entitles users with access to the Edge Threat Management support team via phone and email.

To learn more about this service, visit the Support page. Live Support includes Configuration Backup.

Support

The **Support Portal** button opens our ticketing system, which is also reachable at Support Portal. You can also email us directly to support@arista.com which will automatically create a support ticket.

Please include as much information as possible when filing a ticket, and please remember that a Live Support subscription is required to work directly with the support team. We will be unable to escalate free customer issues to the engineering team.

Support Information lists information that helps Edge Threat Management Support verify your license and identify your box:

- **UID**: The *unique identifier number* for your NG Firewall. This number is generated during the install process.
- **Build**: The exact build version your NG Firewall is running.

Related Topics

- Configuration Backup

Reference Material

Contents

11.1   Event Definitions

All event data is stored in the Mail messages in a relational database. As Arista and applications process traffic they create Event objects that add and modify content in the database. Each event has it's own class/object with certain fields that modify the database in a certain way.

The list below shows the classes used in the event logging and the attributes of each event object. These can be used to add alerts in Reports or for other event handling within Arista.

SpamLogEvent

These events are created by Spam Blocker and update the Database Schema table when an email is scanned.

Attribute Name	Type	Description getAction
action	SpamMessageAction	The action *getClass*
class	Class	The class name *getClientAddr*
clientAddr	InetAddress	The client address *getClientPort*
clientPort	int	The client port getMessageId
messageId	Long	The message ID *getPartitionTablePostfix getReceiver*
receiver	String	The receiver *getScore*
score	float	The score *getSender*
sender	String	The sender *getServerAddr*

Attribute Name	Type	Description getAction
serverAddr	InetAddress	The server address *getServerPort*
serverPort	int	The server port *getSmtpMessageEvent*
smtpMessageEvent	SmtpMessageEvent	The parent SMTP message event isSpam
isSpam	boolean	True if spam, false otherwise getSubject
subject	String	The subject getTag getTestsString
testsString	String	The tests string from the spam engine *getTimeStamp*
timeStamp	Timestamp	The timestamp *getVendorName*
vendorName	String	The application name

SpamSmtpTarpitEvent

These events are created by Spam Blocker and inserted to the Database Schema table when a session is tarpitted.

Attribute Name	Type	Description getIPAddr
IPAddr	InetAddress	The IP address *getClass*
class	Class	The class name *getHostname*
hostname	String	The host name *getPartitionTablePostfix getSessionEvent*
sessionEvent	SessionEvent	The session event *getSessionId*
sessionId	Long	The session ID *getTag getTimeStamp*
timeStamp	Timestamp	The time stamp *getVendorName*
vendorName	String	The application name

PrioritizeEvent

These events are created by the Bandwidth ControlDatabase Schema and update the table when a session is prioritized.

Attribute Name	Type	Description getClass
class	Class	The class name *getPartitionTablePostfix getPriority*
priority	int	The priority *getRuleId*

Attribute Name	Type	Description getClass
ruleId	int	The rule ID ***getSessionEvent***
sessionEvent	SessionEvent	The session event ***getTag*** ***getTimeStamp***
timeStamp	Timestamp	The timestamp

VirusFtpEvent

These events are created by Virus Blocker and update the Database Schema table when Virus Blocker scans an FTP transfer.

Attribute Name	Type	Description getAppName
appName	String	The name of the application ***getClass***
class	Class	The class name ***getClean***
clean	boolean	True if clean, false otherwise ***getPartitionTablePostfix*** ***getSessionEvent***
sessionEvent	SessionEvent	The session event ***getTag*** ***getTimeStamp***
timeStamp	Timestamp	The timestamp ***getUri***
uri	String	The URI ***getVirusName***
virusName	String	The virus name, if not clean

VirusHttpEvent

These events are created by Virus Blocker and update the Database Schema table when Virus Blocker scans an HTTP transfer.

Attribute Name	Type	Description getAppName
appName	String	The name of the application ***getClass***
class	Class	The class name ***getClean***
clean	boolean	True if clean, false otherwise ***getPartitionTablePostfix*** ***getRequestLine***
requestLine	RequestLine	The request line ***getSessionEvent***
sessionEvent	SessionEvent	The session event ***getTag*** ***getTimeStamp***
timeStamp	Timestamp	The timestamp ***getVirusName***
virusName	String	The virus name, if not clean

VirusSmtpEvent

These events are created by Virus Blocker and update the Database Schema table when Virus Blocker scans an email.

Attribute Name	Type	Description getAction
action	String	The action ***getAppName***
appName	String	The name of the application ***getClass***
class	Class	The class name ***getClean***
clean	boolean	True if clean, false otherwise ***getMessageId***
messageId	Long	The message ID ***getPartitionTablePostfix getTag getTimeStamp***
timeStamp	Timestamp	The timestamp ***getVirusName***
virusName	String	The virus name, if not clean

FirewallEvent

These events are created by Firewall and update the Database Schema table when a firewall rule matches a session.

Attribute Name	Type	Description getBlocked
blocked	boolean	True if blocked, false otherwise ***getClass***
class	Class	The class name ***getFlagged***
flagged	boolean	True if flagged, false otherwise ***getPartitionTablePostfix getRuleId***
ruleId	long	The rule ID ***getSessionId***
sessionId	Long	The session ***ID getTag getTimeStamp***
timeStamp	Timestamp	The timestamp

11.2   Day of Week Matcher

A "Day of Week" matcher syntax describes days of the week.

A Day of Week Matcher can be any of the following syntax:

Name	Example	Description
Any Matcher	"any"	matches all days of the week
Single Day (English name)	"tuesday"	matches Tuesday only
Single Day (Digit 1-7)	"1"	matches Sunday only

Name	Example	Description
List of Time of Day Matchers	"monday,2,wednesday"	matches Monday, Tuesday, and Wednesday

11.3   Group Matcher

Group Matcher syntax describes an user or set of users. This can be used, for example, in Policy Manager or Bandwidth Control rules to match against certain traffic.

Group Matcher can be any of the following:

Name	Example	Description
Any Matcher	[any]	matches all groups
None Matcher	[none]	matches no groups
Groupname	mygroup	matches the "mygroup" group
Glob Matcher	m*p	matches the "mygroup" group
List of Group Matchers	mygroup1,mygroup2	matches "mygroup1" and "mygroup2"

11.4   Glob Matcher

A Glob is a common way to match strings of characters against rules. An Arista glob is similar to the syntax commonly used on Microsoft OSs to match filenames (example: "rm *.exe").

A glob matcher has two special characters: "*" means 0 or more of any characters (excluding return charater) and "?" means exactly 1 of any character (excluding return character).

Example	String	Description
String	XYZ	matches "XYZ" but NOT "xYZ" and NOT "XYZZ"
String with *	X*Z	matches "XZ" and "XYZ" and "XYYZ" and "XyyyabcZ" but NOT "xYZ" and NOT "XYZA"
String with *	X*Z*	matches "XZ" and "XYZ" and "XYYZ" and "XyyyabcZ" and "XYZA" but NOT "xYZ"
String with ?	X?Z	matches "XYZ" and "XyZ" but NOT match "XZ" or "XYYZ"
List of Globs	X,Z	matches "X" and "Z" but NOT match "Y" or "X,Z"

Globs are often used in rules like URL rules and filename rules to match various strings. The left and rights side are implicitly anchored. If you wish to match if a string contains the match you will need to use "*foo*".

For those familiar with regular expression you can derive the glob equivalent by doing the following:

- replace "." with "\." to escape the special meaning of "." in regular expressions.
- replace "?" with "." to match any character.
- replace "*" with ".*" to match zero or more characters.

> **Note:**
> - "*" matches all values except null/unset.
> - "" matches null and nothing else.

- All glob matching is case insensitive for domains but case sensitive for all other matches.

11.5   Int Matcher

Int Matcher syntax describes a integer or set of integers.

This can be used, for example, in firewall or Policy manager rules to match against certain traffic destination ports.

Port Matcher can be any of the following:

Name	Example	Description
Any	any	matches all
Single	80	matches that single integer
Greater Than	>1234	matches all values greater than 1234
Less Than	<1234	matches all values less than 1234
Range	1024-65535	matches all values within the range (inclusive)
List of Int Matchers	80,443,8080-8088	matches all 80, 443, and 8080 through 8088

Note: Floating point numbers are also allowed and applicable in some cases. (example: ">2.5")

11.6   IP Matcher

IP Matcher syntax is a that describes an IP address or set of IP addresses. This can be used, for example, in Firewall or Policy Manager rules to match against certain traffic.

IP Matcher can be any of the following:

Name	Example	Description
Any Matcher	any	matches all addresses
Single IP	1.2.3.4	matches the single IP address
Range of IPs	1.2.3.4 - 1.2.3.100	matches all the IPs in the range
CIDR range	192.168.1.0/24	matches all the IPs in that subnet
List of IP Matchers	1.2.3.4, 1.2.3.5, 1.2.3.10, 1.2.3.15	matches all the IPs in the list and in that range

11.7   User Matcher

User Matcher syntax describes an user or set of users.

This can be used, for example, in Policy manager or Bandwidth Application rules to match against certain traffic.

User Matcher can be any of the following:

Name	Example	Description
Any Authenticated User	[authenticated]	matches all identified or authenticated users (excluding null)

Name	Example	Description
Unauthenticated User	[unauthenticated]	matches all unidentified or unauthenticated users (including null)
Username	myuser	matches the "myuser" user
Glob Matcher	m*r	matches the "myuser" user
List of User Matchers	myuser1,myuser2	matches "myuser1" and "myuser2"

11.8 URL Matcher

The URL Matcher Syntax describes all or part of a website.

Example	Matches	Does not Match
example.com	http://example.com/, http://www.example.com/, http://example.com/foo	http://example.net
example.com/bar	http://example.com/bar/test.html, http://www.example.com/bar	http://example.com/foo
porn	http://pornsite.com/	http://foobar.com
example???.com/	http://example123.com	http://example1.com
example.com/foo	http://example.com/foo, http://abc.example.com/foobar	http://example.com/

URL Matchers use globs which are describe more in depth in the Glob Matcher.

Important notes:

- The left side of the rule is anchored with the regular expression "^([a-zA-Z_0-9-]*\.)*". "foo.com" will match only "foo.com" and "abc.foo.com" but not "afoo.com".
- The right side of the rule is anchored with with the regular expression ".*$". "foo.com" will match "foo.com/test.html" because it is actually "foo.com.*$". "foo.com/bar" is "foo.com/bar.*$" which will match "foo.com/bar/baz" and "foo.com/bar2". Also "foo" becomes "foo.*" which will match "foobar.com" and "foo.com".
- "http://" and "https://" are stripped from the rule.
- URIs are case-sensitive, but domains are not. The URL Matcher is case sensitive, but domains are converted to lowercase before evaluation because they should not be case sensitive. Any part of the matcher that should match against the domain should be lower case in the rule.
- "www." is automatically stripped from the rule. This is to prevent the frequent misconfiguration of users adding a block rule for something like "www.pornsite.com" which blocks "www.pornsite.com" but **not** just "pornsite.com." If you truly desire to only match www.pornsite.com and not pornsite.com then use "*www.pornsite.com" because the "*" will match zero or more characters.
- Similarly "*." is stripped from the rule for the same reason as above. If you truly want all subdomains but not the main domain matched, you can accomplish this by doing "*?.foo.com".

11.9 Port Forward Troubleshooting Guide

Port forwards can be tricky. Below is a series of suggestions about getting port forwards to work.

1. Read the Port Forwarding FAQs.
2. Verify that the destination host on the inside is using the Untangle as its default gateway. If not the reply packets won't find their way back to Arista.
3. Verify that the destination service is reachable from the **inside** on the IP and port specified in your port forward rule.
4. Test your (TCP) port forward using 'telnet.' In windows you can run **start** > **run** and then you can type telnet *1.2.3.4 123* where *1.2.3.4* is your external IP and *123* is the port your port forward rule matches. If it connects and hangs then the port forward is working. If it fails to connect then your port forward is not working.
5. Test your rule from the outside. Port forwarding back inside the network has extra complications. First verify that it works from the outside.
6. Verify there is a session shown in Reports > Network > Port Forwarded Sessions
7. Verify that Arista can connect to the final destination. Use the *Connection Test* in *Troubleshooting'* or open the console on Arista and type 'telnet 192.168.1.10 123' where 192.168.1.10 is the internal server you are forwarding to and 123 is the port. If it connects then Arista can reach the server. If it fails to connect Arista can't reach the server and the port forward will probably not function until this part is working.
8. For testing, turn off the Firewall and Captive Portal applications if you have them installed. Port forwarded sessions will not connect if they are blocked by an application. If you have many policies, verify which policy is processing the session and make sure you disable the correct apps.
9. Simplify your port forward rule. Remove extra qualifiers and make it contain as few as possible. For example specify just what port to forward and "Destined Local" and then which server to forward it to. If that works then add the extra qualifiers back one at a time testing each time.
10. If you are port forwarding port *443* (HTTPS), try moving Untangle administration to another port so port *443* is available to be forwarded.
11. Remove any *Source Address* and *Source Interface* qualifiers - 99% of the time these are misused.
12. For advanced users, use tcpdump or the *Packet Test* in troubleshooting to debug and watch the packets. To test with tcpdump run these commands: *tcpdump -i eth0 -n "port 123"* and *tcpdump -i eth1 -n "port 123"* - assuming eth0 is your outside interface and *eth1* is your inside interface.
13. Still not working? Post a **screenshot** of your port forward rule to the forums along with **the results from the above tests** and ask for help.

11.10   Database Schema

Contents

- Database Tables

  - Configuration backup events
  - HTTP events
  - Intrusion prevention events
  - SMPT tarpit events
  - IPsec user events
  - IPsec VPN events
  - IPsec tunnel stats
  - HTTP query events
  - Admin logins
  - Sessions
  - Session minutes
  - Quotas
  - Host table updates
  - Device table updates
  - User table updates

- Alerts
- Settings changes
- Web cache stats
- Server events
- Interface stat events
- Mail messages
- Mail addresses
- FTP events
- Tunnel VPN events
- Tunnel VPN stats
- WAN failover test events
- WAN failover action events
- Directory connector login events
- Captive portal user events
- OpenVPN stats
- OpenVPN events

Database Tables

Configuration backup events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
success	Success	boolean	The result of the backup (true if the backup succeeded, false otherwise)
description	Text detail of the event	text	Text detail of the event
destination	Destination	text	The location of the backup
event_id	Event ID	bigint	The unique event ID

HTTP events

Column Name	Human Name	Type	Description
request_id	Request ID	bigint	The HTTP request ID
time_stamp	Timestamp	timestamp without time zone	The time of the event
session_id	Session ID	bigint	The session
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
c_client_addr	Client-side Client Address	inet	The client-side client IP address
s_client_addr	Server-side Client Address	inet	The server-side client IP address
c_server_addr	Client-side Server Address	inet	The client-side server IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
c_client_port	Client-side Client Port	integer	The client-side client port

Column Name	Human Name	Type	Description
s_client_port	Server-side Client Port	integer	The server-side client port
c_server_port	Client-side Server Port	integer	The client-side server port
s_server_port	Server-side Server Port	integer	The server-side server port
client_country	Client Country	text	The client Country
client_latitude	Client Latitude	real	The client Latitude
client_longitude	Client Longitude	real	The client Longitude
server_country	Server Country	text	The server Country
server_latitude	Server Latitude	real	The server Latitude
server_longitude	Server Longitude	real	The server Longitude
policy_id	Policy ID	smallint	The policy
username	Username	text	The username associated with this session
hostname	Hostname	text	The hostname of the local address
method	Method	character(1)	The HTTP method
uri	URI	text	The HTTP URI
host	Host	text	The HTTP host
domain	Domain	text	The HTTP domain (shortened host)
referer	Referer	text	The Referer URL
c2s_content_length	Client-to-server Content Length	bigint	The client-to-server content length
s2c_content_length	Server-to-client Content Length	bigint	The server-to-client content length
s2c_content_type	Server-to-client Content Type	text	The server-to-client content type
s2c_content_filename	Server-to-client Content Disposition Filename	text	The server-to-client content disposition filename
ad_blocker_cookie_ident	Ad Blocker Cookie	text	This name of cookie blocked by Ad Blocker
ad_blocker_action	Ad Blocker Action	character(1)	This action of Ad Blocker on this request
web_filter_reason	Reason for action (Web Filter)	character(1)	This reason Web Filter blocked/flagged this request
web_filter_category_id	Web Category (Web Filter)	smallint	This numeric category according to Web Filter
web_filter_rule_id	Web Rule (Web Filter)	smallint	This numeric rule according to Web Filter
web_filter_blocked	Blocked (Web Filter)	boolean	If Web Filter blocked this request
web_filter_flagged	Flagged (Web Filter)	boolean	If Web Filter flagged this request
virus_blocker_lite_clean	Virus Blocker Lite Clean	boolean	The cleanliness of the file according to Virus Blocker Lite
virus_blocker_lite_name	Virus Blocker Lite Name	text	The name of the malware according to Virus Blocker Lite

Column Name	Human Name	Type	Description
virus_blocker_clean	Virus Blocker Clean	boolean	The cleanliness of the file according to Virus Blocker
virus_blocker_name	Virus Blocker Name	text	The name of the malware according to Virus Blocker
threat_prevention_blocked	Threat Prevention Blocked	boolean	If Threat Prevention blocked this request
threat_prevention_flagged	Threat Prevention Flagged	boolean	If Threat Prevention flagged this request
threat_prevention_rule_id	Threat Prevention Rule Id	integer	This numeric rule according to Threat Prevention
threat_prevention_reputation	Threat Prevention Reputation	smallint	This numeric threat reputation
threat_prevention_categories	Threat Prevention Categories	integer	This bitmask of threat categories

Intrusion prevention events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
sig_id	Signature ID	bigint	This ID of the rule
gen_id	Grouping ID	bigint	The grouping ID for the rule, The gen_id + sig_id specify the rule's unique identifier
class_id	Classtype ID	bigint	The numeric ID for the classtype
source_addr	Source Address	inet	The source IP address of the packet
source_port	Source Port	integer	The source port of the packet (if applicable)
dest_addr	Destination Address	inet	The destination IP address of the packet
dest_port	Destination Port	integer	The destination port of the packet (if applicable)
protocol	Protocol	integer	The protocol of the packet
blocked	Blocked	boolean	If the packet was blocked/dropped
category	Category	text	The application specific grouping for the signature
classtype	Classtype	text	The generalized threat signature grouping (unrelated to gen_id)
msg	Message	text	The "title" or "description" of the signature
rid	Rule ID	text	The rule id
rule_id	Rule ID	text	The rule id

SMPT Tarpit Events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
ipaddr	Client Address	inet	The client IP address
hostname	Hostname	text	The hostname of the local address
policy_id	Policy ID	bigint	The policy

Column Name	Human Name	Type	Description
vendor_name	Vendor Name	character varying(255)	The "vendor name" of the app that logged the event
event_id	Event ID	bigint	The unique event ID

IPsec User Events

Column Name	Human Name	Type	Description
event_id	Event ID	bigint	The unique event ID
time_stamp	Timestamp	timestamp without time zone	The time of the event
connect_stamp	Connect Time	timestamp without time zone	The time the connection started
goodbye_stamp	End Time	timestamp without time zone	The time the connection ended
client_address	Client Address	text	The remote IP address of the client
client_protocol	Client Protocol	text	The protocol the client used to connect
client_username	Client Username	text	The username of the client
net_process	Net Process	text	The PID of the PPP process for L2TP connections or the connection ID for Xauth connections
net_interface	Net Interface	text	The PPP interface for L2TP connections or the client interface for Xauth connections
elapsed_time	Elapsed Time	text	The total time the client was connected
rx_bytes	Bytes Received	bigint	The number of bytes received from the client in this connection
tx_bytes	Bytes Sent	bigint	The number of bytes sent to the client in this connection

IPsec VPN Events

Column Name	Human Name	Type	Description
event_id	Event ID	bigint	The unique event ID
time_stamp	Timestamp	timestamp without time zone	The time of the event
local_address	Local Address	text	The local address of the tunnel
remote_address	Remote Address	text	The remote address of the tunnel
tunnel_description	Tunnel Description	text	The description of the tunnel
event_type	Event Type	text	The type of the event (CONNECT,DISCONNECT)

IPsec Tunnel Stats

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
tunnel_name	Tunnel Name	text	The name of the IPsec tunnel
in_bytes	In Bytes	bigint	The number of bytes received during this time frame

Column Name	Human Name	Type	Description
out_bytes	Out Bytes	bigint	The number of bytes transmitted during this time frame
event_id	Event ID	bigint	The unique event ID

HTTP Query Events

Column Name	Human Name	Type	Description
event_id	Event ID	bigint	The unique event ID
time_stamp	Timestamp	timestamp without time zone	The time of the event
session_id	Session ID	bigint	The session
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
c_client_addr	Client-side Client Address	inet	The client-side client IP address
s_client_addr	Server-side Client Address	inet	The server-side client IP address
c_server_addr	Client-side Server Address	inet	The client-side server IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
c_client_port	Client-side Client Port	integer	The client-side client port
s_client_port	Server-side Client Port	integer	The server-side client port
c_server_port	Client-side Server Port	integer	The client-side server port
s_server_port	Server-side Server Port	integer	The server-side server port
policy_id	Policy ID	bigint	The policy
username	Username	text	The username associated with this session
hostname	Hostname	text	The hostname of the local address
request_id	Request ID	bigint	The HTTP request ID
method	Method	character(1)	The HTTP method
uri	URI	text	The HTTP URI
term	Search Term	text	The search term
host	Host	text	The HTTP host
c2s_content_length	Client-to-server Content Length	bigint	The client-to-server content length
s2c_content_length	Server-to-client Content Length	bigint	The server-to-client content length
s2c_content_type	Server-to-client Content Type	text	The server-to-client content type
blocked	Blocked	boolean	If Web Filter blocked this search term
flagged	Flagged	boolean	If Web Filter flagged this search term

Admin Logins

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
login	Login	text	The login name
local	Local	boolean	True if it is a login attempt through a local process
client_addr	Client Address	inet	The client IP address
succeeded	Succeeded	boolean	True if the login succeeded, false otherwise
reason	Reason	character(1)	The reason for the login (if applicable)

Sessions

Column Name	Human Name	Type	Description
session_id	Session ID	bigint	The session
time_stamp	Timestamp	timestamp without time zone	The time of the event
end_time	End Time	timestamp without time zone	The time the session ended
bypassed	Bypassed	boolean	True if the session was bypassed, false otherwise
entitled	Entitled	boolean	True if the session is entitled to premium functionality
protocol	Protocol	smallint	The IP protocol of session
icmp_type	ICMP Type	smallint	The ICMP type of session if ICMP
hostname	Hostname	text	The hostname of the local address
username	Username	text	The username associated with this session
policy_id	Policy ID	smallint	The policy
policy_rule_id	Policy Rule ID	smallint	The ID of the matching policy rule (0 means none)
local_addr	Local Address	inet	The IP address of the local participant
remote_addr	Remote Address	inet	The IP address of the remote participant
c_client_addr	Client-side Client Address	inet	The client-side client IP address
c_server_addr	Client-side Server Address	inet	The client-side server IP address
c_server_port	Client-side Server Port	integer	The client-side server port
c_client_port	Client-side Client Port	integer	The client-side client port
s_client_addr	Server-side Client Address	inet	The server-side client IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
s_server_port	Server-side Server Port	integer	The server-side server port
s_client_port	Server-side Client Port	integer	The server-side client port
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
client_country	Client Country	text	The client Country

Column Name	Human Name	Type	Description
client_latitude	Client Latitude	real	The client Latitude
client_longitude	Client Longitude	real	The client Longitude
server_country	Server Country	text	The server Country
server_latitude	Server Latitude	real	The server Latitude
server_longitude	Server Longitude	real	The server Longitude
c2p_bytes	From-Client Bytes	bigint	The number of bytes the client sent to Arista (client-to-pipeline)
p2c_bytes	To-Client Bytes	bigint	The number of bytesArista sent to client (pipeline-to-client)
s2p_bytes	From-Server Bytes	bigint	The number of bytes the server sent to Arista (client-to-pipeline)
p2s_bytes	To-Server Bytes	bigint	The number of bytes Arista sent to server (pipeline-to-client)
filter_prefix	Filter Block	text	The network filter that blocked the connection (filter,shield,invalid)
firewall_blocked	Firewall Blocked	boolean	True if Firewall blocked the session, false otherwise
firewall_flagged	Firewall Flagged	boolean	True if Firewall flagged the session, false otherwise
firewall_rule_index	Firewall Rule ID	integer	The matching rule in Firewall (if any)
threat_prevention_blocked	Threat Prevention Blocked	boolean	If Threat Prevention blocked
threat_prevention_flagged	Threat Prevention Flagged	boolean	If Threat Prevention flagged
threat_prevention_reason	Threat Prevention Reason	character(1)	Threat Prevention reason
threat_prevention_rule_id	Threat Prevention Rule Id	integer	Numeric rule id of Threat Prevention
threat_prevention_client_reputation	Threat Prevention Client Reputation	smallint	Numeric client reputation of Threat Prevention
threat_prevention_client_categories	Threat Prevention Client Categories	integer	Bitmask client categories of Threat Prevention
threat_prevention_server_reputation	Threat Prevention Server Reputation	smallint	Numeric server reputation of Threat Prevention
threat_prevention_server_categories	Threat Prevention Server Categories	integer	Bitmask server categories of Threat Prevention
application_control_lite_protocol	Application Control Lite Protocol	text	The application protocol according to Application Control Lite
application_control_lite_blocked	Application Control Lite Blocked	boolean	True if Application Control Lite blocked the session
captive_portal_blocked	Captive Portal Blocked	boolean	True if Captive Portal blocked the session
captive_portal_rule_index	Captive Portal Rule ID	integer	The matching rule in Captive Portal (if any)
application_control_application	Application Control Application	text	The application according to Application Control
application_control_protochain	Application Control Protochain	text	The protochain according to Application Control
application_control_category	Application Control Category	text	The category according to Application Control
application_control_blocked	Application Control Blocked	boolean	True if Application Control blocked the session
application_control_flagged	Application Control Flagged	boolean	True if Application Control flagged the session

Column Name	Human Name	Type	Description
application_control_confidence	Application Control Confidence	integer	True if Application Control confidence of this session's identification
application_control_ruleid	Application Control Rule ID	integer	The matching rule in Application Control (if any)
application_control_detail	Application Control Detail	text	The text detail from the Application Control engine
bandwidth_control_priority	Bandwidth Control Priority	integer	The priority given to this session
bandwidth_control_rule	Bandwidth Control Rule ID	integer	The matching rule in Bandwidth Control rule (if any)
ssl_inspector_ruleid	SSL Inspector Rule ID	integer	The matching rule in SSL Inspector rule (if any)
ssl_inspector_status	SSL Inspector Status	text	The status/action of the SSL session (INSPECTED, IGNORED, BLOCKED, UNTRUSTED, ABANDONED)
ssl_inspector_detail	SSL Inspector Detail	text	Additional text detail about the SSL connection (SNI, IP Address)
tags	Tags	text	The tags on this session

Session Minutes

Column Name	Human Name	Type	Description
session_id	Session ID	bigint	The session
time_stamp	Timestamp	timestamp without time zone	The time of the event
c2s_bytes	From-Client Bytes	bigint	The number of bytes the client sent
s2c_bytes	From-Server Bytes	bigint	The number of bytes the server sent
start_time	Start Time	timestamp without time zone	The start time of the session
end_time	End Time	timestamp without time zone	The time the session ended
bypassed	Bypassed	boolean	True if the session was bypassed, false otherwise
entitled	Entitled	boolean	True if the session is entitled to premium functionality
protocol	Protocol	smallint	The IP protocol of session
icmp_type	ICMP Type	smallint	The ICMP type of session if ICMP
hostname	Hostname	text	The hostname of the local address
username	Username	text	The username associated with this session
policy_id	Policy ID	smallint	The policy
policy_rule_id	Policy Rule ID	smallint	The ID of the matching policy rule (0 means none)
local_addr	Local Address	inet	The IP address of the local participant
remote_addr	Remote Address	inet	The IP address of the remote participant
c_client_addr	Client-side Client Address	inet	The client-side client IP address
c_server_addr	Client-side Server Address	inet	The client-side server IP address
c_server_port	Client-side Server Port	integer	The client-side server port
c_client_port	Client-side Client Port	integer	The client-side client port

Column Name	Human Name	Type	Description
s_client_addr	Server-side Client Address	inet	The server-side client IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
s_server_port	Server-side Server Port	integer	The server-side server port
s_client_port	Server-side Client Port	integer	The server-side client port
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
client_country	Client Country	text	The client Country
client_latitude	Client Latitude	real	The client Latitude
client_longitude	Client Longitude	real	The client Longitude
server_country	Server Country	text	The server Country
server_latitude	Server Latitude	real	The server Latitude
server_longitude	Server Longitude	real	The server Longitude
filter_prefix	Filter Block	text	The network filter that blocked the connection (filter,shield,invalid)
firewall_blocked	Firewall Blocked	boolean	True if Firewall blocked the session, false otherwise
firewall_flagged	Firewall Flagged	boolean	True if Firewall flagged the session, false otherwise
firewall_rule_index	Firewall Rule ID	integer	The matching rule in Firewall (if any)
threat_prevention_blocked	Threat Prevention Blocked	boolean	If Threat Prevention blocked
threat_prevention_flagged	Threat Prevention Flagged	boolean	If Threat Prevention flagged
threat_prevention_reason	Threat Prevention Reason	character(1)	Threat Prevention reason
threat_prevention_rule_id	Threat Prevention Rule Id	integer	Numeric rule id of Threat Prevention
threat_prevention_client_reputation	Threat Prevention Client Reputation	smallint	Numeric client reputation of Threat Prevention
threat_prevention_client_categories	Threat Prevention Client Categories	integer	Bitmask client categories of Threat Prevention
threat_prevention_server_reputation	Threat Prevention Server Reputation	smallint	Numeric server reputation of Threat Prevention
threat_prevention_server_categories	Threat Prevention Server Categories	integer	Bitmask server categories of Threat Prevention
application_control_lite_protocol	Application Control Lite Protocol	text	The application protocol according to Application Control Lite
application_control_lite_blocked	Application Control Lite Blocked	boolean	True if Application Control Lite blocked the session
captive_portal_blocked	Captive Portal Blocked	boolean	True if Captive Portal blocked the session
captive_portal_rule_index	Captive Portal Rule ID	integer	The matching rule in Captive Portal (if any)
application_control_application	Application Control Application	text	The application according to Application Control
application_control_protochain	Application Control Protochain	text	The protochain according to Application Control
application_control_category	Application Control Category	text	The category according to Application Control

Column Name	Human Name	Type	Description
application_control_blocked	Application Control Blocked	boolean	True if Application Control blocked the session
application_control_flagged	Application Control Flagged	boolean	True if Application Control flagged the session
application_control_confidence	Application Control Confidence	integer	True if Application Control confidence of this session's identification
application_control_ruleid	Application Control Rule ID	integer	The matching rule in Application Control (if any)
application_control_detail	Application Control Detail	text	The text detail from the Application Control engine
bandwidth_control_priority	Bandwidth Control Priority	integer	The priority given to this session
bandwidth_control_rule	Bandwidth Control Rule ID	integer	The matching rule in Bandwidth Control rule (if any)
ssl_inspector_ruleid	SSL Inspector Rule ID	integer	The matching rule in SSL Inspector rule (if any)
ssl_inspector_status	SSL Inspector Status	text	The status/action of the SSL session (INSPECTED, IGNORED, BLOCKED, UNTRUSTED, ABANDONED)
ssl_inspector_detail	SSL Inspector Detail	text	Additional text detail about the SSL connection (SNI, IP Address)
tags	Tags	text	The tags on this session

Quotas

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
entity	Entity	text	The IP entity given the quota (address/ username)
action	Action	integer	The action (1=Quota Given, 2=Quota Exceeded)
size	Size	bigint	The size of the quota
reason	Reason	text	The reason for the action

Host Table Updates

Column Name	Human Name	Type	Description
address	Address	inet	The IP address of the host
key	Key	text	The key being updated
value	Value	text	The new value for the key
old_value	Old Value	text	The old value for the key
time_stamp	Timestamp	timestamp without time zone	The time of the event

Device Table Updates

Column Name	Human Name	Type	Description
mac_address	MAC Address	text	The MAC address of the device
key	Key	text	The key being updated
value	Value	text	The new value for the key
old_value	Old Value	text	The old value for the key

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event

User Table Updates

Column Name	Human Name	Type	Description
username	Username	text	The username
key	Key	text	The key being updated
value	Value	text	The new value for the key
old_value	Old Value	text	The old value for the key
time_stamp	Timestamp	timestamp without time zone	The time of the event

Alerts

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
description	Text detail of the event	text	The description from the alert rule.
summary_text	Summary Text	text	The summary text of the alert
json	JSON Text	text	The summary JSON representation of the event causing the alert

Settings Changes

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
settings_file	Settings File	text	The name of the file changed
username	Username	text	The username logged in at the time of the change
hostname	Hostname	text	The remote hostname

Web cache stats

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
hits	Hits	bigint	The number of cache hits during this time frame
misses	Misses	bigint	The number of cache misses during this time frame
bypasses	Bypasses	bigint	The number of cache user bypasses during this time frame
systems	System bypasses	bigint	The number of cache system bypasses during this time frame
hit_bytes	Hit Bytes	bigint	The number of bytes saved from cache hits
miss_bytes	Miss Bytes	bigint	The number of bytes not saved from cache misses

Column Name	Human Name	Type	Description
event_id	Event ID	bigint	The unique event ID

Server events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
load_1	CPU load (1-min)	numeric(6,2)	The 1-minute CPU load
load_5	CPU load (5-min)	numeric(6,2)	The 5-minute CPU load
load_15	CPU load (15-min)	numeric(6,2)	The 15-minute CPU load
cpu_user	CPU User Utilization	numeric(6,3)	The user CPU percent utilization
cpu_system	CPU System Utilization	numeric(6,3)	The system CPU percent utilization
mem_total	Total Memory	bigint	The total bytes of memory
mem_free	Memory Free	bigint	The number of free bytes of memory
disk_total	Disk Size	bigint	The total disk size in bytes
disk_free	Disk Free	bigint	The free disk space in bytes
swap_total	Swap Size	bigint	The total swap size in bytes
swap_free	Swap Free	bigint	The free disk swap in bytes
active_hosts	Active Hosts	integer	The number of active hosts

Interface stat events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
interface_id	Interface ID	integer	The interface ID
rx_rate	Rx Rate	double precision	The RX rate (bytes/s)
rx_bytes	Bytes Received	bigint	The number of bytes received from the client in this connection
tx_rate	Tx Rate	double precision	The TX rate (bytes/s)
tx_bytes	Bytes Sent	bigint	The number of bytes sent to the client in this connection

Mail messages

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
session_id	Session ID	bigint	The session
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
c_client_addr	Client-side Client Address	inet	The client-side client IP address
s_client_addr	Server-side Client Address	inet	The server-side client IP address
c_server_addr	Client-side Server Address	inet	The client-side server IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
c_client_port	Client-side Client Port	integer	The client-side client port

Column Name	Human Name	Type	Description
s_client_port	Server-side Client Port	integer	The server-side client port
c_server_port	Client-side Server Port	integer	The client-side server port
s_server_port	Server-side Server Port	integer	The server-side server port
policy_id	Policy ID	bigint	The policy
username	Username	text	The username associated with this session
msg_id	Message ID	bigint	The message ID
subject	Subject	text	The email subject
hostname	Hostname	text	The hostname of the local address
event_id	Event ID	bigint	The unique event ID
sender	Sender	text	The address of the sender
receiver	Receiver	text	The address of the receiver
virus_blocker_lite_clean	Virus Blocker Lite Clean	boolean	The cleanliness of the file according to Virus Blocker Lite
virus_blocker_lite_name	Virus Blocker Lite Name	text	The name of the malware according to Virus Blocker Lite
virus_blocker_clean	Virus Blocker Clean	boolean	The cleanliness of the file according to Virus Blocker
virus_blocker_name	Virus Blocker Name	text	The name of the malware according to Virus Blocker
spam_blocker_lite_score	Spam Blocker Lite Score	real	The score of the email according to Spam Blocker Lite
spam_blocker_lite_is_spam	Spam Blocker Lite Spam	boolean	The spam status of the email according to Spam Blocker Lite
spam_blocker_lite_tests_string	Spam Blocker Lite Tests	text	The tess results for Spam Blocker Lite
spam_blocker_lite_action	Spam Blocker Lite Action	character(1)	The action taken by Spam Blocker Lite
spam_blocker_score	Spam Blocker Score	real	The score of the email according to Spam Blocker
spam_blocker_is_spam	Spam Blocker Spam	boolean	The spam status of the email according to Spam Blocker
spam_blocker_tests_string	Spam Blocker Tests	text	The tess results for Spam Blocker
spam_blocker_action	Spam Blocker Action	character(1)	The action taken by Spam Blocker
phish_blocker_score	Phish Blocker Score	real	The score of the email according to Phish Blocker
phish_blocker_is_spam	Phish Blocker Phish	boolean	The phish status of the email according to Phish Blocker
phish_blocker_tests_string	Phish Blocker Tests	text	The tess results for Phish Blocker
phish_blocker_action	Phish Blocker Action	character(1)	The action taken by Phish Blocker

Mail addresses

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
session_id	Session ID	bigint	The session
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
c_client_addr	Client-side Client Address	inet	The client-side client IP address
s_client_addr	Server-side Client Address	inet	The server-side client IP address

c_server_addr	Client-side Server Address	inet	The client-side server IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
c_client_port	Client-side Client Port	integer	The client-side client port
s_client_port	Server-side Client Port	integer	The server-side client port
c_server_port	Client-side Server Port	integer	The client-side server port
s_server_port	Server-side Server Port	integer	The server-side server port
policy_id	Policy ID	bigint	The policy
username	Username	text	The username associated with this session
msg_id	Message ID	bigint	The message ID
subject	Subject	text	The email subject
addr	Address	text	The address of this event
addr_name	Address Name	text	The name for this address
addr_kind	Address Kind	character(1)	The type for this address (F=From, T=To, C=CC, G=Envelope From, B=Envelope To, X=Unknown)
hostname	Hostname	text	The hostname of the local address
event_id	Event ID	bigint	The unique event ID
sender	Sender	text	The address of the sender
virus_blocker_lite_clean	Virus Blocker Lite Clean	boolean	The cleanliness of the file according to Virus Blocker Lite
virus_blocker_lite_name	Virus Blocker Lite Name	text	The name of the malware according to Virus Blocker Lite
virus_blocker_clean	Virus Blocker Clean	boolean	The cleanliness of the file according to Virus Blocker
virus_blocker_name	Virus Blocker Name	text	The name of the malware according to Virus Blocker
spam_blocker_lite_score	Spam Blocker Lite Score	real	The score of the email according to Spam Blocker Lite
spam_blocker_lite_is_spam	Spam Blocker Lite Spam	boolean	The spam status of the email according to Spam Blocker Lite
spam_blocker_lite_action	Spam Blocker Lite Action	character(1)	The action taken by Spam Blocker Lite
spam_blocker_lite_tests_string	Spam Blocker Lite Tests	text	The tess results for Spam Blocker Lite
spam_blocker_score	Spam Blocker Score	real	The score of the email according to Spam Blocker
spam_blocker_is_spam	Spam Blocker Spam	boolean	The spam status of the email according to Spam Blocker
spam_blocker_action	Spam Blocker Action	character(1)	The action taken by Spam Blocker
spam_blocker_tests_string	Spam Blocker Tests	text	The tess results for Spam Blocker
phish_blocker_score	Phish Blocker Score	real	The score of the email according to Phish Blocker
phish_blocker_is_spam	Phish Blocker Phish	boolean	The phish status of the email according to Phish Blocker
phish_blocker_tests_string	Phish Blocker Tests	text	The tess results for Phish Blocker
phish_blocker_action	Phish Blocker Action	character(1)	The action taken by Phish Blocker

FTP events

Column Name	Human Name	Type	Description
event_id	Event ID	bigint	The unique event ID

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
session_id	Session ID	bigint	The session
client_intf	Client Interface	smallint	The client interface
server_intf	Server Interface	smallint	The server interface
c_client_addr	Client-side Client Address	inet	The client-side client IP address
s_client_addr	Server-side Client Address	inet	The server-side client IP address
c_server_addr	Client-side Server Address	inet	The client-side server IP address
s_server_addr	Server-side Server Address	inet	The server-side server IP address
policy_id	Policy ID	bigint	The policy
username	Username	text	The username associated with this session
hostname	Hostname	text	The hostname of the local address
request_id	Request ID	bigint	The FTP request ID
method	Method	character(1)	The FTP method
uri	URI	text	The FTP URI
virus_blocker_lite_clean	Virus Blocker Lite Clean	boolean	The cleanliness of the file according to Virus Blocker Lite
virus_blocker_lite_name	Virus Blocker Lite Name	text	The name of the malware according to Virus Blocker Lite
virus_blocker_clean	Virus Blocker Clean	boolean	The cleanliness of the file according to Virus Blocker
virus_blocker_name	Virus Blocker Name	text	The name of the malware according to Virus Blocker

Tunnel VPN events

Column Name	Human Name	Type	Description
event_id	Event ID	bigint	The unique event ID
time_stamp	Timestamp	timestamp without time zone	The time of the event
tunnel_name	Tunnel Name	text	The name the tunnel
server_address	Server IP Address	text	The address of the remote server
local_address	Local Address	text	The local address assigned the client
event_type	Event Type	text	The type of the event (CONNECT,DISCONNECT)

Tunnel VPN stats

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
tunnel_name	Tunnel Name	text	The name of the Tunnel VPN tunnel
in_bytes	In Bytes	bigint	The number of bytes received during this time frame

Column Name	Human Name	Type	Description
out_bytes	Out Bytes	bigint	The number of bytes transmitted during this time frame
event_id	Event ID	bigint	The unique event ID

WAN failover test events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
interface_id	Interface ID	integer	This interface ID
name	Interface Name	text	This name of the interface
description	Text detail of the event	text	The description from the test rule
success	Success	boolean	The result of the test (true if the test succeeded, false otherwise)
event_id	Event ID	bigint	The unique event ID

WAN failover action events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
interface_id	Interface ID	integer	This interface ID
action	Action	text	This action (CONNECTED,DISCONNECTED)
os_name	Interface O/S Name	text	This O/S name of the interface
name	Interface Name	text	This name of the interface
event_id	Event ID	bigint	The unique event ID

Directory connector login events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
login_name	Login Name	text	The login name
domain	Domain	text	The AD domain
type	Type	text	The type of event (I=Login,U=Update,O=Logout)
client_addr	Client Address	inet	The client IP address
login_type	Login Type	text	The login type

Captive portal user events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
policy_id	Policy ID	bigint	The policy
event_id	Event ID	bigint	The unique event ID

Column Name	Human Name	Type	Description
login_name	Login Name	text	The login username
event_info	Event Type	text	The type of event (LOGIN, FAILED, TIMEOUT, INACTIVE, USER_LOGOUT, ADMIN_LOGOUT)
auth_type	Authorization Type	text	The authorization type for this event
client_addr	Client Address	text	The remote IP address of the client

OpenVPN stats

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
start_time	Start Time	timestamp without time zone	The time the OpenVPN session started
end_time	End Time	timestamp without time zone	The time the OpenVPN session ended
rx_bytes	Bytes Received	bigint	The total bytes received from the client during this session
tx_bytes	Bytes Sent	bigint	The total bytes sent to the client during this session
remote_address	Remote Address	inet	The remote IP address of the client
pool_address	Pool Address	inet	The pool IP address of the client
remote_port	Remote Port	integer	The remote port of the client
client_name	Client Name	text	The name of the client
event_id	Event ID	bigint	The unique event ID

OpenVPN events

Column Name	Human Name	Type	Description
time_stamp	Timestamp	timestamp without time zone	The time of the event
remote_address	Remote Address	inet	The remote IP address of the client
pool_address	Pool Address	inet	The pool IP address of the client
client_name	Client Name	text	The name of the client
type	Type	text	The type of the event (CONNECT,DISCONNECT)

11.11   Rules

Rules are used frequently in Arista and many other firewalls. Rules are very powerful, but can sometimes be difficult to configure.

This documentation describes how rules work and gives some basic examples and some common mistakes to avoid. Many applications use rules like firewall, Captive Portal, Bandwidth Application control, etc. All of these rules essentially share the same logic.

Basics

Rules are configured by the user to categorize and act upon traffic. For example, firewall uses rules to determine whether to block or pass traffic. Bandwidth Application control uses rules to determine how to prioritize a session.

Rules are evaluated in order from top to bottom against sessions (not packets!). If a rule matches then the action from that rule is performed and no more rules are evaluated. If no matching rule is found the behavior is defined by the application, which is usually doing nothing.

This is similar to other firewalls. Arista rules are **quick** rules which means the first match is always used. Unlike some other firewalls, Arista evaluates against sessions, not packets. This means that the **Source Address** will be the initiator of the session and the **Destination Address** will be the server address the client is connecting to, and the same is true for **Source Port** and **Destination Port**.

Each rule has several properties:

- An enable checkbox
- A name/description
- A set of conditions
- An action or set of actions

Rule Anatomy

The enable checkbox determines if the rule is evaluated. If the enable checkbox is not checked, the rule is simply skipped.

The description is for the user to document what the rule does. It is *highly* suggested to give the rule a meaningful name. Trying to troubleshoot a set of rules all named "[no description]" can be extremely difficult.

The set of conditions is the description of the traffic that should match the rule. If *all* of the conditions are true for the given session then the rules matches. This is discussed in more details in the next section.

The action or set of actions configures what action is performed if the rule matches. This is dependent on the application. For example, in firewall it determines whether to block or pass the session and additionally if it should be flagged.

Conditions

Conditions define which sessions will match the rule. If and only if all of the conditions match, the rule is considered a match. Conditions can also be inverted by selecting "is NOT" in the dropdown, effectively inverting when it matches. *Destination Port* "is NOT" "*80*" matches on all ports except port *80*.

Lets take a simple example: We want to block TCP traffic to port *80* on a server. There is a web service running on that server that you don't want to allow access to. First, create a rule and check the **enable** checkbox and give it a reasonably descriptive name like "blocking TCP port *80* to serverX." Now you will need to add some conditions that match only the traffic you want to block. So in this example will will add:

- **Protocol** is TCP
- **Destination Address** is *1.2.3.4* (The IP address of the server)
- **Destination Port** is *80*

Finally, set the action to block and flag, and click **Done** and **Apply**. Since all conditions must be true this rule will block TCP traffic to *1.2.3.4* port *80* only and nothing else.

There are many conditions available to carefully define precise sets of traffic. The following table defines the list of conditions.

Each condition has several properties:

- Name - The *Name* of the condition.
- Syntax - The accepted *Syntax* of the condition. If editing through the UI, some conditions have custom editors to help you craft conditions, others are just a text field.

- Availability - The *Availability* is the availability of that matcher at various times. Some conditions, like *Destination Address*, are always known and can always be used to match. Other conditions, like *Application Control: Application*, are only known after the session is created and some data flows and Application Control is able to identify the application. These type conditions are 'deep session' conditions because they are not able to be evaluated at session creation time, only after some data flows. As such they are not available in rules that are evaluated at session creation time, like firewall and Captive Portal.
- Reliability - The *Reliability* is a "reliability" of that condition. *True* means it is 100% reliable. *False* means it is 99% or less reliable. For example, some conditions, like *Destination Port*, are always known and thus matching on *Destination Port* is 100% reliable. Other conditions, like *Client Hostname*, only match if the hostname for the client hostname is known. Hostname is determined through many ways, including DNS and DHCP. If all of these methods fail then it is entirely possible that the hostname of a server is "foo" but Arista has not been able to determine this and as such a *Client Hostname* is "foo" condition will not match. This column is there purely as a warning that users in these cases must be aware of when conditions might deliver false negatives.

Table 19: Condition List

Name	Syntax	Function	Availability	Reliability
Destination Address	IP Matcher	Matches if value matches the Destination/Server Address of the session. (after NAT/port forwarding)	all	True
Destination Port	Int Matcher	Matches if value matches the Destination/Server Port of the session. (after NAT/port forwarding)	all	True
Destination Interface	checkboxes	Matches if value matches the Destination/Server Interface of the session. (after routing/port forwarding)	all	True
Destined Local	boolean	Matches if the session is destined to one of Arista's IP addresses.	all	True
Source Address	IP Matcher	Matches if value matches the Source/Client Address of the session. (before NAT/port forwarding)	all	True
Source Port	NG Firewall Rule Syntax	Matches if value matches the Source/Client Port of the session. (before NAT/port forwarding)	hidden	True
Source Interface	checkboxes	Matches if value matches the Source/Client Interface of the session. (before routing/port forwarding)	all	True
Protocol	checkboxes	Matches if value matches the Protocol of the session.	all	True
Tagged	Glob Matcher	Matches if session is tagged with matching tag	all	True

Name	Syntax	Function	Availability	Reliability
Client Tagged	Glob Matcher	Matches if client of the session is tagged with matching tag	all	True
Server Tagged	Glob Matcher	Matches if server of the session is tagged with matching tag	all	True
Username	User Matcher	Matches if value matches the username associated with the Client IP in the Host table.	all	False
Host Hostname	Glob Matcher	Matches if value matches the hostname associated with the Local Address in the Host table.	all	False
Client Hostname	Glob Matcher	Matches if value matches the hostname associated with the Client IP in the Host table.	all	False
Server Hostname	Glob Matcher	Matches if value matches the hostname associated with the Server IP in the Host table.	all	False
Client MAC Address	Glob Matcher	Matches if value matches the MAC address associated with the Client IP in the ARP table.	all	False
Server MAC Address	Glob Matcher	Matches if value matches the MAC address associated with the Server IP in the ARP table.	all	False
Client MAC Vendor	Glob Matcher	Matches if value matches the device manufacturer. This is identified using the Client IP's MAC address in the ARP table and OUI Lookup.	all	False
Server MAC Vendor	Glob Matcher	Matches if value matches the device manufacturer. This is identified using the Server IP's MAC address in the ARP table and OUI Lookup.	all	False
Host has no Quota	boolean	Matches if the local address has no quota	all	True
User has no Quota	boolean	Matches if the username has no quota	all	True
Client has no Quota	boolean	Matches if the client has no quota	hidden	True
Server has no Quota	boolean	Matches if the server has no quota	hidden	True
Host has Exceeded Quota	boolean	Matches if the local address has exceeded their quota	all	True

Name	Syntax	Function	Availability	Reliability
User has Exceeded Quota	boolean	Matches if the user has exceeded their quota	all	True
Client has Exceeded Quota	boolean	Matches if the client has exceeded their quota	hidden	True
Server has Exceeded Quota	boolean	Matches if the server has exceeded their quota	hidden	True
Host Quota Attainment	Int Matcher	Matches the local address quota used to quota size ratio (">1" means over quota, ">.5" means over 50% used, etc)	all	True
User Quota Attainment	Int Matcher	Matches the username quota used to quota size ratio (">1" means over quota, ">.5" means over 50% used, etc)	all	True
Client Quota Attainment	Int Matcher	Matches the client quota used to quota size ratio (">1" means over quota, ">.5" means over 50% used, etc)	hidden	True
Server Quota Attainment	Int Matcher	Matches the server quota used to quota size ratio (">1" means over quota, ">.5" means over 50% used. etc)	hidden	True
HTTP: Hostname	Glob Matcher	Matches if the value matches the hostname specified in the HTTP session	deep sessions	False
HTTP: Referrer	Glob Matcher	Matches if the value matches the referrer string specified in the HTTP header	deep sessions	False
HTTP: URI	Glob Matcher	Matches if the value matches the latest URI specified in the HTTP session	deep sessions	False
HTTP: URL	URL Matcher	Matches if the value matches the latest URL (hostname +URI) specified in the HTTP session	deep sessions	False
HTTP: Content Type	Glob Matcher	Matches the content-type of the latest content in the HTTP session	deep sessions	False
HTTP: Content Length	Int Matcher	Matches if Int Matcher matches the content length specified in the latest HTTP response	deep sessions	False

Name	Syntax	Function	Availability	Reliability
HTTP: Request Method	Glob Matcher	Matches if the value matches the HTTP request method of the requested URL. Standard request methods include GET, POST, HEAD, OPTIONS, PUT, DELETE, TRACE, and CONNECT.	deep sessions	False
HTTP: Request File Path	Glob Matcher	Matches if the value matches the entire file path of the requested URL. This is everything including and after the first slash character following the host name or address. (e.g. /some/location/ mypage.html)	deep sessions	False
HTTP: Request File Name	Glob Matcher	Matches if the value matches the file name of the requested URL. This is everything after the final slash character of the request. (e.g. mypage.html)	deep sessions	False
HTTP: Request File Extension	Glob Matcher	Matches if the value matches the file extension of the requested URL. This is everything after following the dot after the final slash character of the request. (e.g. html)	deep sessions	False
HTTP: Response Content Type	Glob Matcher	Matches if the value matches the content or MIME type reported in the server response.	deep sessions	False
HTTP: Response File Name	Glob Matcher	Matches if the value matches the file name returned in the server response.	deep sessions	False
HTTP: Response File Extension	Glob Matcher	Matches if the value matches the file extension returned in the server response. This is everything after (but not including) the final dot of the file name.	deep sessions	False
HTTP: Client User Agent	Glob Matcher	Matches if the value matches the User Agent string for the client in the Host table	all	False
Application Control: Application	Glob Matcher	Matches if the value matches the Application determined by Application Control	deep sessions	False

Name	Syntax	Function	Availability	Reliability
Application Control: Category	Glob Matcher	Matches if the value matches the Category of the Application determined by Application Control	deep sessions	False
Application Control: Protochain	Glob Matcher	Matches if the value matches the Protochain determined by Application Control	deep sessions	False
Application Control: Detail	Glob Matcher	Matches if the value matches the Detail of the Application determined by Application Control	deep sessions	False
Application Control: Confidence	Int Matcher	Matches if Int Matcher matches the confidence rating determined by Application Control	deep sessions	False
Application Control: Productivity	Int Matcher	Matches if Int Matcher matches the productivity rating determined by Application Control	deep sessions	False
Application Control: Risk	Int Matcher	Matches if Int Matcher matches the risk rating determined by Application Control	deep sessions	False
Application Control Lite: Signature	Glob Matcher	Matches if the value matches the Signature determined by Application Control Lite	deep sessions	False
Application Control Lite: Category	Glob Matcher	Matches if the value matches the Category of the Signature determined by Application Control Lite	deep sessions	False
Application Control Lite: Description	Glob Matcher	Matches if the value matches the Description of the Signature determined by Application Control Lite	deep sessions	False
Web Filter: Category	Glob Matcher	Matches if the value matches the Category determined by Web Filter	deep sessions	False
Web Filter: Category Description	Glob Matcher	Matches if the value matches the Description of the Category determined by Web Filter	deep sessions	False
Web Filter: Site is Flagged	boolean	Matches if the latest request in this session was flagged by Web Filter	deep sessions	False

Name	Syntax	Function	Availability	Reliability
Directory Connector: User in Group	Group Matcher	Matches if the username associated with the client in the Host Table is in the specified group(s)	all	False
SSL Inspector: SNI Host Name	Glob Matcher	Matches if the value matches the Server Name Indication (SNI) host name included by the client in the initial session request.	deep sessions	False
SSL Inspector: Certificate Subject	Glob Matcher	Matches if the value matches the Subject DN in the SSL certificate received from the external server.	all	False
SSL Inspector: Certificate Issuer	Glob Matcher	Matches if the value matches the Issuer DN in the SSL certificate received from the external server.	all	False
Time of Day	Time and Date Formatting	Matches times of day.	all	True
Day of Week	Day of Week Matcher	Matches days of the week.	all	True
Remote Host Country	Glob Matcher	Matches if value matches the country associated with the remote IP.	all	False
Client Country	Glob Matcher	Matches if value matches the country associated with the Client IP.	all	False
Server Country	Glob Matcher	Matches if value matches the country associated with the Server IP.	all	False

Order

As discussed above, the order of rules is very important. Often users want to do very complex tasks that can be difficult or impossible with a single rule. In these cases, multiple rules is useful. For example, assume you want to use Firewall to block all traffic to a server (*1.2.3.4*) except port *80* traffic or from the admin IP (*192.168.1.100*) to RDP (port *3389*).

Doing this in one rule would be difficult (actually impossible in this case), but it is quite easy using several rules: First, create a rule to just block all traffic to the server IP, *1.2.3.4*. Above that rule create a rule that passes all traffic from the admin IP, *192.168.1.100*, to the server IP where *Destination Port* == "*3389*", and another rule that passes all traffic to *1.2.3.4* with *Destination Port* == "*80*." This set of three rules does exactly what we describe above and effects no other traffic.

	Rule Id	Enable	Description	Conditions	Block	Flag	Edit	Delete
+	100001	☑	Allow traffic to 1.2.3.4 when Source Address = 192.168.1.100 and Destination Port = 3389	Source Address ➡ **192.168.1.100** • Destination Port ➡ **3389** • Destination Address ➡ **1.2.3.4**	☐	☐	✏	🗑
+	100002	☑	Allow traffic to 1.2.3.4 when Destination Port = 80	Destination Port ➡ **80** • Destination Address ➡ **1.2.3.4**	☐	☐	✏	🗑
+	100003	☑	Block traffic to 1.2.3.4	Destination Address ➡ **1.2.3.4**	☑	☑	✏	🗑

Common Mistakes

This is a list of common mistakes to avoid.

- **Conditions must ALL match**

    In order for a rule to match, ALL conditions must match. In some cases you may want to add a rule that just passes all traffic to and from an IP. Users sometimes will create a pass rule with two conditions:

    - *Destination Address* is "*1.2.3.4*"
    - *Source Address* is "*1.2.3.4*"

    This rule will never match anything because it will only match traffic where the destination address is *1.2.3.4* AND the source address is *1.2.3.4* (a session destined to itself). If you want to match when the destination address is *1.2.3.4* OR the source address is *1.2.3.4*, then you must create two separate rules, one for the destination address and one for the source address.

- **Rule Order**

    Occasionally, you will add a rule and not understand why it is not working as intended. Often we find that there was a rule above that was matching the session before the desired rule. For example, just adding a rule to the bottom often isn't sufficient, you must find the appropriate place in your ruleset to place the new rule.

    To do so simply logically evaluate the rules in your head starting at the top to find the appropriate place for your rule. Also use the event log to run several tests to see how the session was handled in the event log and adjust your ruleset accordingly.

11.11.1   NG Firewall Rule Syntax

About NG Firewall Rule Syntax

Throughout the NG Firewall Server Administrative Interface, Administrators must enter information about their network and web locations. In some cases the values entered can be exact, and in others the text entered indicates a range of values.

The following describe common syntaxes to describe IPs, ports, strings, URLs, etc. The Policy Manager documentation describes which syntax is used for which fields.

IP Matcher

IP Matcher syntax is a that describes an IP address or set of IP addresses. This can be used, for example, in firewall or Policy Manager rules to match against certain traffic.

IP Matcher can be any of the following:

Name	Example	Description
Any Matcher	any	matches all addresses
Single IP	*1.2.3.4*	matches the single IP address
Range of IPs	*1.2.3.4*-*1.2.3.100*	matches all the IPs in the range
CIDR range	*192.168.1.0/24*	matches all the IPs in that subnet
List of IP Matchers	*1.2.3.4*, *1.2.3.5*, *1.2.3.10* - *1.2.3.15*	matches all the IPs in the list and in that range

Port Matcher

Int Matcher syntax is a that describes a integer or set of integers. This can be used, for example, in firewall or Policy Manager rules to match against certain traffic destination ports.

Port Matcher can be any of the following:

Name	Example	Description
Any	any	matches all
Single	***80***	matches that single integer
Greater Than	***>1234***	matches all values greater than ***1234***
Less Than	***<1234***	matches all values less than ***1234***
Range	***1024-65535***	matches all values within the range (inclusive)
List of Int Matchers	***80***, ***443***, ***8080-8088***	matches all ***80***, ***443***, and ***8080*** through ***8088***

> 📝 **Note:** Floating point numbers are also allowed and apply in some cases. (example: ">2.5")

URL Matcher

The URL Matcher Syntax describes all or part of a website.

Example	Matches	Does not Match
example.com	http://example.com/, http://www.example.com/, http://example.com/foo	http://example.net
example.com/bar	http://example.com/bar/test.html, http://www.example.com/bar	http://example.com/foo
porn	http://pornsite.com/	http://foobar.com
example???.com	http://example123.com	http://example1.com
example.com/foo	http://example.com/foo, http://abc.example.com/foobar	http://example.com/

URL Matchers use globs which are describe more in depth in the Glob Matcher.

> 📝 **Note:**
> - The left side of the rule is anchored with the regular expression "^([a-zA-Z_0-9-]*\.)*". "foo.com" will match only "foo.com" and "abc.foo.com" but not "afoo.com".
> - The right side of the rule is anchored with with the regular expression ".*$". "foo.com" will match "foo.com/test.html" because it is actually "foo.com.*$". "foo.com/bar" is "foo.com/bar.*$" which will match "foo.com/bar/baz" and "foo.com/bar2". Also "foo" becomes "foo.*" which will match "foobar.com" and "foo.com"
> - "http://" and "https://" are stripped from the rule.
> - URIs are case-sensitive, but domains are not. The URL Matcher is case sensitive, but domains are converted to lowercase before evaluation because they should not be case sensitive. Any part of the matcher that should match against the domain should be lower case in the rule.

- "www." is automatically stripped from the rule. This is to prevent the frequent misconfiguration of users adding a block rule for something like "www.pornsite.com" which blocks "www.pornsite.com" but **not** just "pornsite.com." If you truly desire to only match www.pornsite.com and not pornsite.com then use "*www.pornsite.com" because the "*" will match zero or more characters.
- Similarly "*." is stripped from the rule for the same reason as above. If you truly want all subdomains but not the main domain matched, you can accomplish this by doing "*?.foo.com"

User Matcher

User Matcher syntax describes an user or set of users. This can be used, for example, in Policy Manager or Quota rules to match against certain traffic.

User Matcher can be any of the following:

Name	Example	Description
Any Authenticated User	[authenticated]	matches all identified or authenticated users (excluding null)
Unauthenticated User	[unauthenticated]	matches all unidentified or unauthenticated users (including null)
Username	myuser	matches the "myuser" user
Glob Matcher	m*r	matches the "myuser" user
List of User Matchers	myuser1,myuser2	matches "myuser1" and "myuser2"

Group Matcher

Group Matcher syntax is a that describes an user or set of users. This can be used, for example, in Policy Manager or Bandwidth Application Control rules to match against certain traffic.

Group Matcher can be any of the following:

Name	Example	Description
Any Matcher	[any]	matches all groups
None Matcher	[none]	matches no groups
Groupname	mygroup	matches the "mygroup" group
Glob Matcher	m*p	matches the "mygroup" group
List of Group Matchers	mygroup1,mygroup2	matches "mygroup1" and "mygroup2"

Glob Matcher

A Glob is a common way to match strings of characters against rules. An Arista glob is similar to the syntax commonly used on Microsoft OSs to match filenames (example: "rm *.exe").

A glob matcher has two special characters: "*" means 0 or more of any characters (excluding return charater) and "?" means exactly 1 of any character (excluding return character).

Example	String	Description
String	XYZ	matches "XYZ" but NOT "xYZ" and NOT "XYZZ"
String with *	X*Z	matches "XZ" and "XYZ" and "XYYZ" and "XyyyabcZ" but NOT "xYZ" and NOT "XYZA"
String with *	X*Z*	matches "XZ" and "XYZ" and "XYYZ" and "XyyyabcZ" and "XYZA" but NOT "xYZ"
String with ?	X?Z	matches "XYZ" and "XyZ" but NOT match "XZ" or "XYYZ"
List of Globs	X,Z	matches "X" and "Z" but NOT match "Y" or "X,Z"

Globs are often used in rules like URL rules and filename rules to match various strings. The left and rights side are implicitly anchored. If you wish to match if a string contains the match you will need to use "*foo*".

For those familiar with regular expression you can derive the glob equivalent by doing the following:

- replace "." with "\." to escape the special meaning of "." in regular expressions.
- replace "?" with "." to match any character.
- replace "*" with ".*" to match zero or more characters.

Note:

- "*" matches all values except null/unset.
- "" matches null and nothing else.
- All glob matching is case insensitive for domains but case sensitive for all other matches.

Time of Day Matcher

Day of Week Matcher

Day of Week Matcher

A "Day of Week" matcher is a syntax used to describe days of the week.

A Day of Week Matcher can be any of the following syntax:

Name	Example	Description
Any Matcher	"any"	matches all days of the week
Single Day (English name)	"tuesday"	matches Tuesday only
Single Day (Digit 1-7)	"1"	matches Sunday only
List of Time of Day Matchers	"monday, 2, wednesday"	matches Monday, Tuesday, and Wednesday

11.12 Time and Date Formatting

Format	Description	Example
d	Day of the month, 2 digits with leading zeros	01 to 31
D	A short textual representation of the day of the week	Mon to Sun
j	Day of the month without leading zeros	1 to 31
l	A full textual representation of the day of the week	Sunday to Saturday
N	ISO-8601 numeric representation of the day of the week	1 (for Monday) through 7 (for Sunday)
S	English ordinal suffix for the day of the month, 2 characters. Works well with j	st, nd, rd or th
w	Numeric representation of the day of the week	0 (for Sunday) to 6 (for Saturday)
z	The day of the year (starting from 0)	0 to 364 (365 in leap years)
W	ISO-8601 week number of year, weeks starting on Monday	01 to 53
F	A full textual representation of a month, such as January or March	January to December
m	Numeric representation of a month, with leading zeros	01 to 12
M	A short textual representation of a month	Jan to Dec
n	Numeric representation of a month, without leading zeros	1 to 12
t	Number of days in the given month	28 to 31
L	Whether it's a leap year	1 if it is a leap year, 0 otherwise.
o	ISO-8601 year number (identical to (Y), but if the ISO week number (W) belongs to the previous or next year, that year is used instead)	Examples: 1998 or 2004
Y	A full numeric representation of a year, 4 digits	Examples: 1999 or 2003
y	A two digit representation of a year	Examples: 99 or 03
a	Lowercase Ante meridiem and Post meridiem	am or pm
A	Uppercase Ante meridiem and Post meridiem	AM or PM
g	12-hour format of an hour without leading zeros	1 to 12

Format	Description	Example
G	24-hour format of an hour without leading zeros	0 to 23
h	12-hour format of an hour with leading zeros	01 to 12
H	24-hour format of an hour with leading zeros	00 to 23
i	Minutes, with leading zeros	00 to 59
s	Seconds, with leading zeros	00 to 59
u	Decimal fraction of a second	Examples:(minimum 1 digit, arbitrary number of digits allowed) 001 (i.e. 0.001s) or 100 (i.e. 0.100s) or 999 (i.e. 0.999s) or 999876543210 (i.e. 0.999876543210s)
O	Difference to Greenwich time (GMT) in hours and minutes	Example: +1030
P	Difference to Greenwich time (GMT) with colon between hours and minutes	Example: -08:00
T	Timezone abbreviation of the machine running the code	Examples: EST, MDT, PDT ...
Z	Timezone offset in seconds (negative if west of UTC, positive if east)	-43200 to 50400
c	ISO 8601 date represented as the local time with an offset to UTC appended.Notes: 1) If unspecified, the month / day defaults to the current month / day, the time defaults to midnight, while the timezone defaults to the browser's timezone. If a time is specified, it must include both hours and minutes. The "T" delimiter, seconds, milliseconds and timezone are optional. 2) The decimal fraction of a second, if specified, must contain at least 1 digit (there is no limit to the maximum number of digits allowed), and may be delimited by either a '.' or a ',' Refer to the examples on the right for the various levels of date-time granularity which are supported, or see http://www.w3.org/TR/NOTE-datetime for more info.	Examples:1991 or 1992-10 or 1993-09-20 or 1994-08-19T16:20+01:00 or 1995-07-18T17:21:28-02:00 or 1996-06-17T18:22:29.98765+03:00 or 1997-05-16T19:23:30,12345-0400 or 1998-04-15T20:24:31.2468Z or 1999-03-14T20:24:32Z or 2000-02-13T21:25:33 2001-01-12 22:26:34
C	An ISO date string as implemented by the native Date object's Date to ISO String method. This outputs the numeric part with *UTC* hour and minute values, and indicates this by appending the `'Z'` timezone identifier.	1962-06-17T09:21:34.125Z

Format	Description	Example
U	Seconds since the Unix Epoch (January 1 1970 00:00:00 GMT)	1193432466 or -2138434463
MS	Microsoft AJAX serialized dates	\/Date(1238606590509)\/ (i.e. UTC milliseconds since epoch) or \/ Date(1238606590509+0800)\/
time	A javascript millisecond timestamp	1350024476440
timestamp	A UNIX timestamp (same as U)	1350024866