FortiWAN Administrator Guide

Version 5.1.2

TABLE OF CONTENTS

Getting Started

Best practice is to follow this work flow for a new deployment.

- Configuration changes are applied to the running configuration as soon as you save them.
- Configuration objects are saved in a configuration management database. You can't change the name of a configuration object after you have initially saved it.
- You can't delete a configuration object that's referenced in another configuration object. For example, you can't delete an address used in a policy.

Step 1: Install the appliance

Prerequisites

- The appliance is installed in a hardware rack or the virtual appliance installed into a VMware environment.

For information on hardware appliances, refer to the FortiWAN hardware manuals.

For information on the virtual appliance, see the FortiWAN-VM Install Guide at https://docs.fortinet.com/fortiwan/hardware.

Step 2: Configure the management interface

Use the management port for administrator access, and also for management traffic (such as SNMP or syslog). If your appliance has a dedicated management port, that's the port you configure as the management interface; otherwise, conventionally use port 4 (VM and FWN30E use Port 3).

Configure the following basic settings to get started so you can access the UI from a remote location (such as your desk):

- **Static route** - Enter the gateway router for the management subnet so you can access the UI from a host on your subnet.

- **IP address** - You typically assign a static IP address for the management interface. The IP address is the host portion of the UI URL. For example, the default IP address for the management interface is `192.168.1.99` and the default URL for the UI is `https://192.168.1.99`.
- **Access** - Services for administrative access. We recommend `HTTPS, SSH, SNMP, PING`.

Prerequisites

- Administrator access. See Administrator access on page 30.
- Know the IP address for the default gateway of the management subnet and the IP address that you plan to assign the management interface.
- Access to the machine room in which a physical appliance has been installed. With physical appliances, you must connect a cable to the management port to get started.
- Have a laptop with an RJ-45 Ethernet network port, a crossover Ethernet cable, and a web browser (a recent version of Chrome, Firefox, or Internet Explorer).
- Configure the laptop Ethernet port with the static IP address `192.168.1.2` and a net mask of `255.255.255.0`. These settings enable you to access the FortiWAN UI as if from the same subnet as the FortiWAN in its factory configuration state.

To connect to the UI

1. Use the crossover cable to connect the laptop Ethernet port to the FortiWAN management port.
2. On your laptop, open the following URL in your browser: `https://192.168.1.99/`

   The system brings up a self-signed security certificate, which it shows to clients whenever they initiate an HTTPS connection.
3. Verify and accept the certificate, and acknowledge any warnings about self-signed certificates. The administrator login page appears.

4. Enter the user name **admin** and no password.

   The dashboard appears.

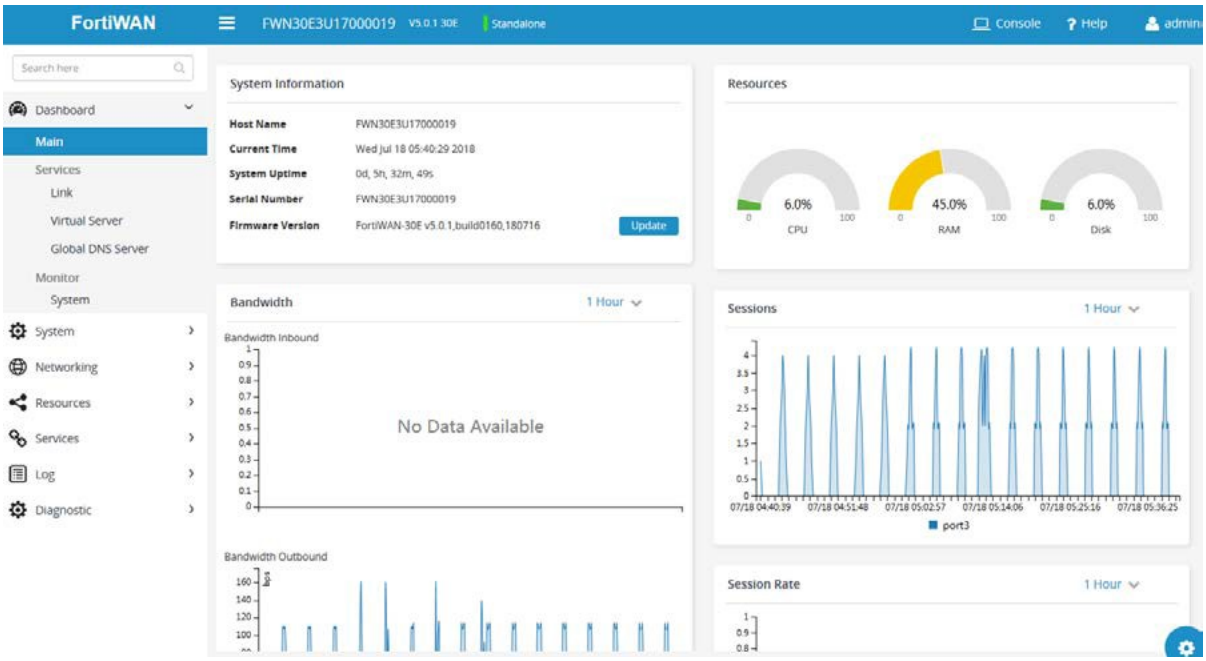

You can also access the UI using CLI commands. See CLI: Connect to the UI on page 235.

Step 3: Configure basic network settings

The system supports network settings for various environments. To get started, you configure the following basic settings:

- **Administrator password** - Change the password for the admin account.
- **System date and time** - We recommend you use NTP to maintain the system time.
- **Network interfaces** - Configure interfaces to receive and forward the network traffic to and from the destination servers.
- **DNS** - A primary and secondary server for system DNS lookups.

Prerequisites

- Know the IP address for the NTP servers your network uses to maintain system time.
- Know the IP addresses that have been provisioned for the traffic interfaces for your FortiWAN deployment.
- Know the IP address for the primary and secondary DNS servers your network uses for DNS resolution.

To change the admin password

- See Admin on page 31 For more information on administrator accounts, see Admin on page 31.

You can also perform this action using CLI commands. See CLI: Change Admin password on page 234.

To configure system time

- See Maintenance on page 39.

You can also perform this action using CLI commands. See CLI: Set System Time on page 236.

To configure network interfaces

- See Interface settings on page 63.

You can also perform this action using CLI commands. See CLI Configure Network Interfaces on page 237.

To configure DNS

- See Basic on page 37. For information on configuring DNS, see System on page 30.

You can also perform this action using CLI commands. See CLI: Configure DNS on page 236.

Step 4: Test connectivity to destination servers

Use `ping` and `traceroute` to test connectivity to destination servers.

To test connectivity from the FortiWAN system to the destination server

Run the following commands from the CLI:

```
execute ping <destination_ip4>
execute traceroute <destination_ipv4>
```

To test connectivity from the destination server to the FortiWAN system

1. Enable ping on the network interface.
2. Use the `ping` and `traceroute` utilities available on the destination server to test connectivity to the FortiWAN network interface IP address.

If you have trouble connecting, see Connectivity issues on page 269.

Step 5: Complete product registration, licensing, and upgrades

Your new FortiWAN appliance comes with a factory image of the operating system (firmware). However, if a new version has been released since factory imaging, you might want to install the newer firmware before continuing the system configuration.

Prerequisites

- Register - Registration is required to log into the Fortinet Customer Service & Support site and download firmware upgrade files. For details, go to

http://kb.fortinet.com/kb/documentLink.do?externalID=12071.

- Check the installed firmware version - In the UI, go to **System > Administrator > Maintenance**, and scroll to the bottom of the page. See Maintenance on page 39.
- Check for upgrades - Major releases include new features, enhancements, and bug fixes. Patch releases can include enhancements and bug fixes.
- Download the release notes at http://docs.fortinet.com/fortiwan/release-information.
- Download firmware upgrades at https://support.fortinet.com/.

Upload your license and new firmware

1. Within the FortiWAN UI, go to **Dashboard > Main**, and within the **License Information** pane, click **Update**.
2. **Upload** the license file.
3. Within the **System Information** pane, click **Update** next to Firmware Version.
4. **Upload** the firmware file.
   - See Updating firmware on page 40.

Step 6: Configure a basic Link load balancing and SNAT rules

A FortiWAN Link load balancing has many custom configuration options. You can leverage the predefined Link, link group, flow-policy configurations to get started.

1. By default, Port1 and Port2 are WAN ports.
2. Configure two SNAT rules for Port1 and Port2.

See Links - underlay or overlay on page 76 and Link group on page 116.

To configure SNAT rules, see Source NAT on page 163.

- Select Port1 as Out Interface.
- Add SNAT rules for Port2 using the same method.

Step 7: Test the deployment

You can test the load balancing deployment by emulating the traffic flow of your planned production deployment.

Figure 1 - Basic network topology

Test basic load balancing

1. Send multiple client requests to the Internet IP address/domain name.
2. Go to **Dashboard > Main** to see the Sessions and throughput counters increment.
3. Go to **Dashboard > Services > Link** to see the link statistics.

Examples

Following are examples of the logs and reports you can use to verify your deployment.

Figure 2 - Dashboard report

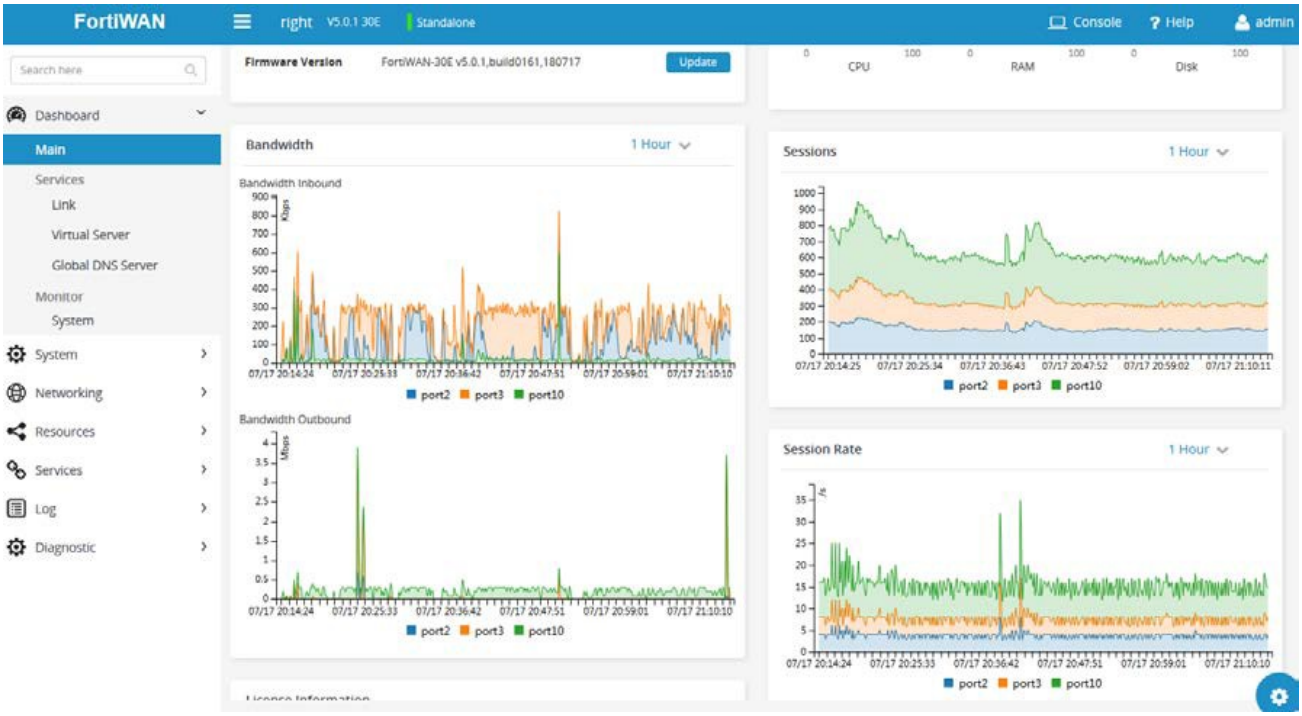

Figure 3 - Link statistics

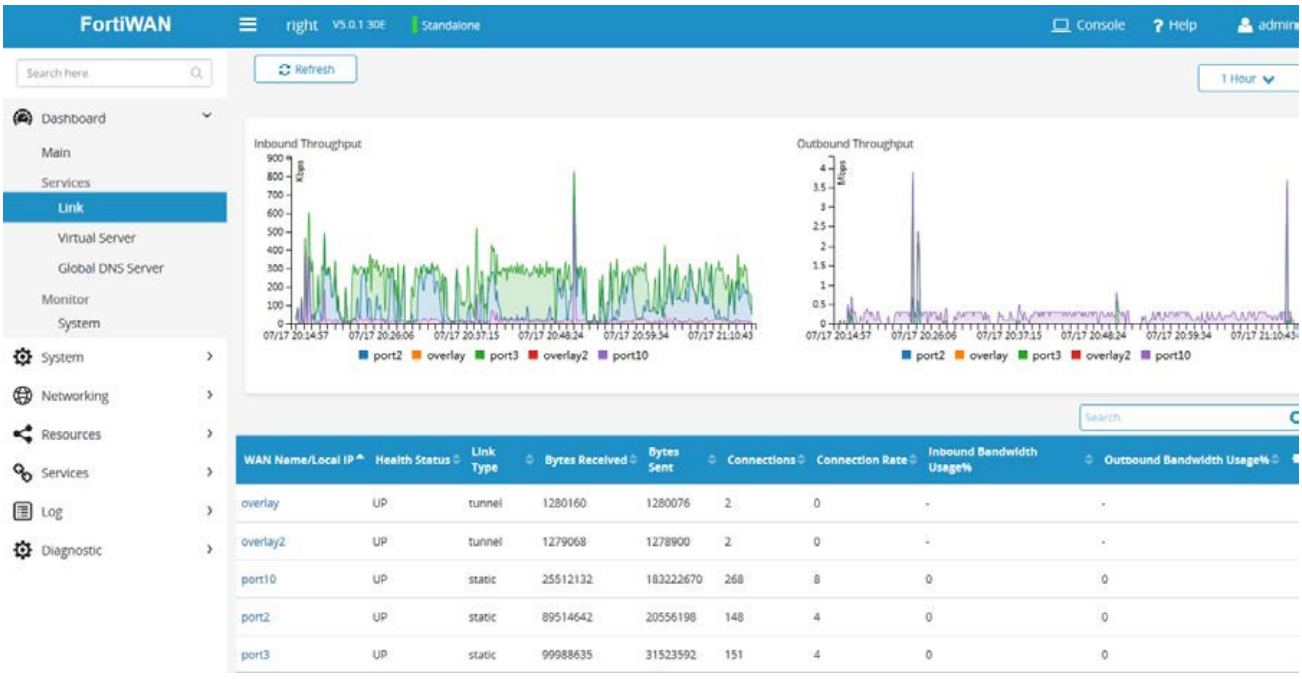

Step 8: Back up the configuration

Once you have tested your basic installation and verified that it functions correctly, create a backup. This “clean” backup is a reference point with many benefits:

- **Troubleshooting** - You can use a diff tool to compare a problematic configuration with this baseline configuration.
- **Restarting** - You can rapidly restore your system to a simple yet working point.
- **Rapid deployment** - You can use the configuration file as a template for other FortiWAN systems. You can use any text editor to edit the plain text configuration file and import it into another FortiWAN system. You should change unique identifiers, such as IP address and other local network settings that differ from one deployment to another.

See Back up and restore on page 42.

What's New

This section highlights the main features of the FortiWan 5.1.2 release.

WebUI enhancements

When system license is invalid or expires, GUI will be redirected to license import page.

LLB health check enhancement

Change the default health check method of llb links to be LLB_HLTHCK_ICMP + LLB_HLTHCK_HOPS; the relationship is OR.

FortiWAN Cloud-Init

FortiWAN Cloud-Init Config Drive is used to pre-configure the FortiWAN VM so that it boots with a valid license and a pre-determined configuration.

FortiWAN UI

Open the FortiWan UI in your browser.

- We only recommend running FortiWAN on Windows 10+, using Chrome or Firefox. Other platforms and browsers have not been tested.

Parts of the UI

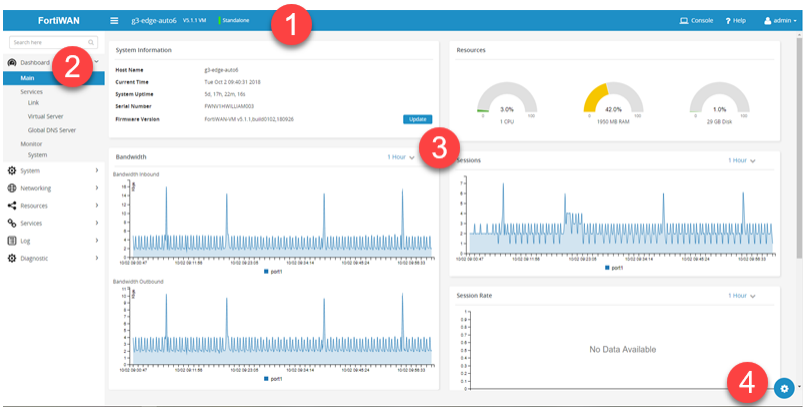

1. **Header** - Shows the version and provides view options.
   - Click the hamburger to toggle the **Menu**.
   - Click **Console** to enter console mode.
   - Click **Help** to open the online help.
   - Click the down arrow next to your user name. See Logout, reboot, reset, and shut down the system on page 19
     - **Logout** - Logs the current user out of the system without shutting it down.
     - **Reboot** - Reboots the operating system.
     - **Reset** - Resets the configuration to the default factory values.
     - **Shutdown** - Shuts down the system. When the system is shut down, it is unavailable to forward traffic.
2. **Menu** - select an item from the list to view information in the Main Window. See Dashboard on page 21.
3. **Main window** - shows information chosen from the left menu.
4. **Dashboard editor** - click to **Add**, **Edit**, **Delete**, or **Reset** dashboards, and to add **widgets** to a dashboard.

Logout, reboot, reset, and shut down the system

You can reboot, reset, logout, or shut down the system from the UI.

- **Logout** - Logs the current user out of the system without shutting it down.
  - On the right-upper corner of the page, click the account icon to expand the drop-down menu, click **Logout**.
- **Reboot** - Reboots the operating system. Do one of the following:
  - On the right-upper corner of the page, click the account icon to expand the drop-down menu, click **Reboot**.
  - From the CLI console, enter:

    ```
    execute reboot
    ```
- **Reset** - Resets the configuration to the default factory values. Do one of the following:
  - On the right-upper corner of the page, click the account icon to expand the drop-dwon menu, click **Reset Configuration**.
  - From the CLI console, enter:

    ```
    execute factoryreset
    ```
- **Shutdown** - Shuts down the system. When the system is shut down, it is unavailable to forward traffic.

  Do one of the following:
  - On the right-upper corner of GUI, click the account icon to expand the drop down menu, then click **Shutdown**.
  - From the CLI console, enter:

    ```
    execute shutdown
    ```

  

  Don't unplug or switch off the FortiWAN appliance without first shutting down the operating system. The shutdown process enables the system to finish writing any buffered data, and to correctly spin down and park the hard disks. Failure to do so could cause data loss and hardware problems.
- **To completely power off**
  - For hardware appliances, press the power button if there is one. Power supplies and switches vary by hardware model. On some, you press the power button; on others, you

flip the switch to either the off (O) or on (I) position.

- For FortiWAN-VM, power off the virtual machine.

Dashboard

The system dashboard appears when you log into the system (or into a virtual domain). It enables you to monitor system-wide health and utilization.

Figure 4 - Main Dashboard UI

Configuring the Dashboard

You can create, modify, or delete a dashboard, or reset all dashboards to default configuration.

Access

- From the Dashboard, go to **Main**, then select the gear icon.

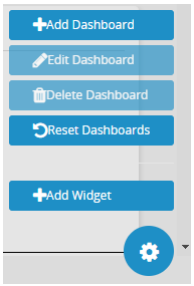

To create a dashboard

1. Click **Add Dashboard**.

2. Enter a name, then click **Save**.

   The new empty dashboard appears in the list.

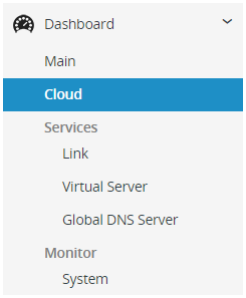

   Add widgets to your dashboard.

To modify a dashboard

Modify a dashboard by adding or removing widgets, or re-arranging widgets.

1. Click **Add Widget**, then the Add Widget window appears.
2. Toggle **ON** the widgets you want.
3. Click **Save**.

   The widgets appear on your dashboard.
4. To rearrange, click in the header area of the widget, then drag to a new location.

   The arrangement is automatically saved.

To delete a dashboard

You can’t delete any default dashboards. You can only delete dashboards you have created.

1. Go to the dashboard you want to delete.
2. Click **Delete Dashboard**, then the delete confirmation window appears.
3. Click **Delete**.

   The dashboard no longer appears.

To reset the dashboards

Resetting the dashboards sets all dashboards back to the default configuration and removes any global VDOM dashboards and widgets.

1. Click **Reset Dashboards**, then the Reset Dashboard window appears.
2. Click **Reset**.

It might take a minute to update the dashboards, depending on your internet connection.

You might need to refresh the page to see the changes.

Adjusting Time Frames in Widgets

Various widget allow you to view data in specific time frames. For example, on the Main dashboard, you can adjust your view of the inbound and outbound time frame. The default time in 1 hour.

1. From the widget header, click the time designation to get the drop down list of available time frames.

2. Select a time frame from the list.

   The graph adjusts accordingly.

   Data does not appear for any time the device has been offline.

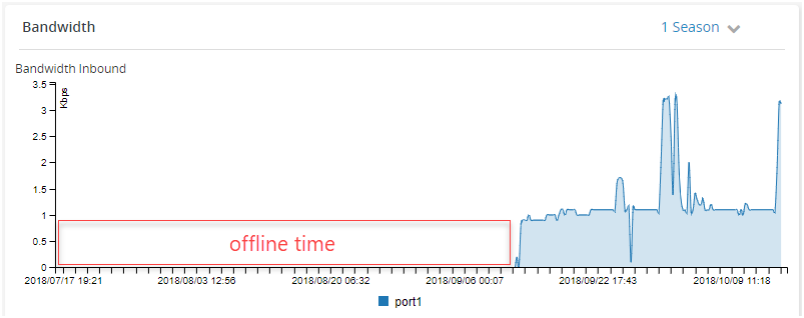

Main

The **Main** tab opens when you select Dashboard from the side menu. This page is divided into five panels:

Porlet	Description
System Information	host name, current time, system up time, serial number, firmware version. Operations: Update firmware, upload license, reboot, shutdown, reset. • Click **Update** to update the FortiWAN firmware.

Porlet	Description
Resources	CPU utilization, Memory utilization, disk utilization, concurrent connections, connections per second, inbound throughput, outbound throughput.
License Information	License status, support contract information, and FortiGuard services version. **Operations**: Upload license, navigate to the support site, or navigate to the FortiGuard services configuration page. • Click **Login** to register your FortiWAN. • Click **Login and Update** to update your FortiWAN license.
Bandwidth (graph)	Shows inbound and outbound traffic.
Sessions (graph)	Shows concurrent connections and connections per second.

Services

The Services dashboard shows information about WAN services.

Link

This page shows information about link load-balancing.

Access

- From the Dashboard, go to **Services > Link**.

Figure 5 - Link UI

1. Click **Refresh** to update the live data.
2. Select a time frame from the list - from 1 hour to 1 year. See Adjusting Time Frames in Widgets on page 23.
3. Shows the live inbound traffic. The legend below the graph corresponds to the port names.
4. Shows the live outbound traffic. The legend below the graph corresponds to the port names.
5. Click a column heading to sort in alpha or reverse alpha order.
6. Click a port name to view details. A new page appears, in which you can zoom in for a more granular view.
7. Search for a port name to filter your view.
8. Select how many entries to view on a page. You can choose 10, 25, or 50.
9. Click **Previous** or **Next** to see more pages. This only works if you have multiple pages.

Virtual Server

This page enables you to check virtual server statistics. See Virtual server settings on page 191.

Access

- From the Dashboard, go to **Services > Virtual Server**.

Figure 6 - Virtual Server UI

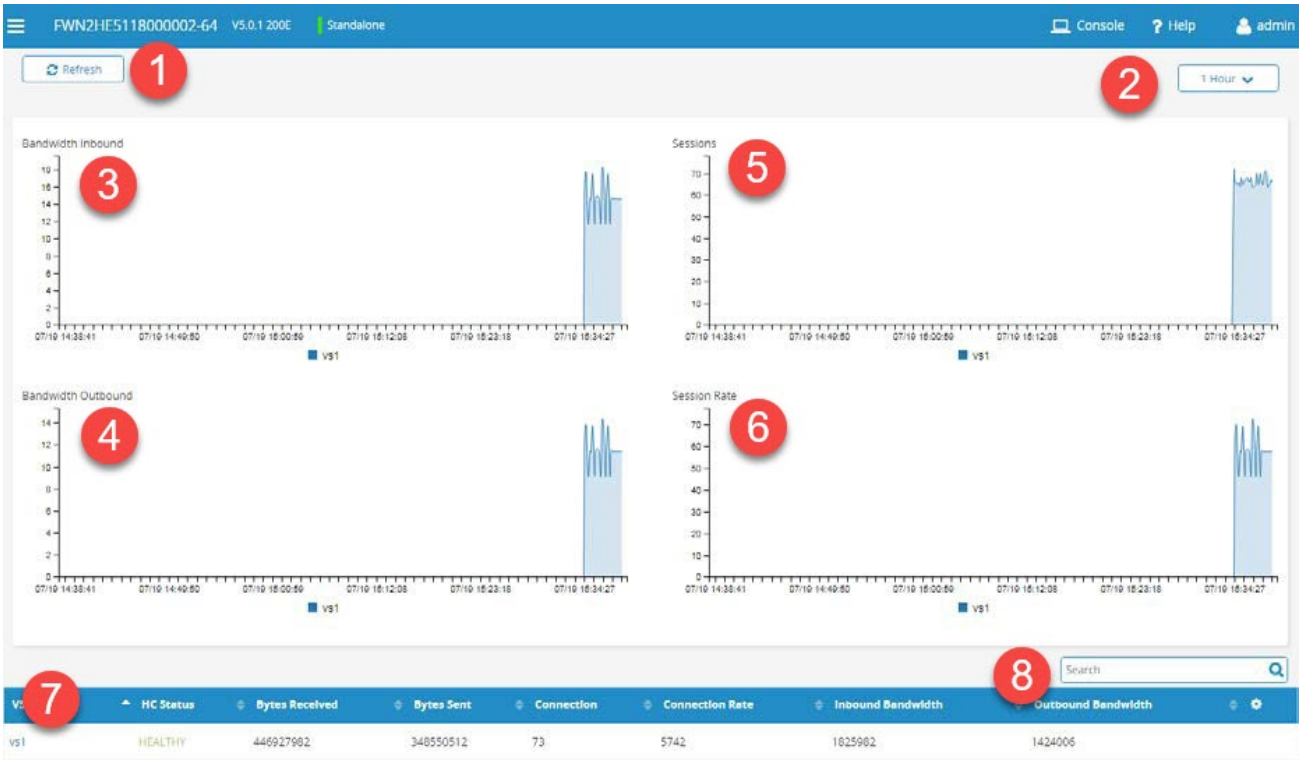

1. Click **Refresh** to update the live data.
2. Select a time frame from the list - from 1 hour to 1 year. See Adjusting Time Frames in Widgets on page 23.
3. Shows the inbound bandwidth.
4. Shows the outbound bandwidth.
5. Shows the number of sessions.
6. Shows the session rate.
7. Click a column heading to sort in alpha or reverse alpha order.

   Click a port name to view details. A new page appears, in which you can zoom in for a more granular view.:
8. Search for a port name to filter your view.

Global DNS Server

This page shows the information about the DNS servers used in global load-balancing.

Access: From the Dashboard, go to **Services > Global DNS Server**.

Figure 7 - Global DNS Server UI

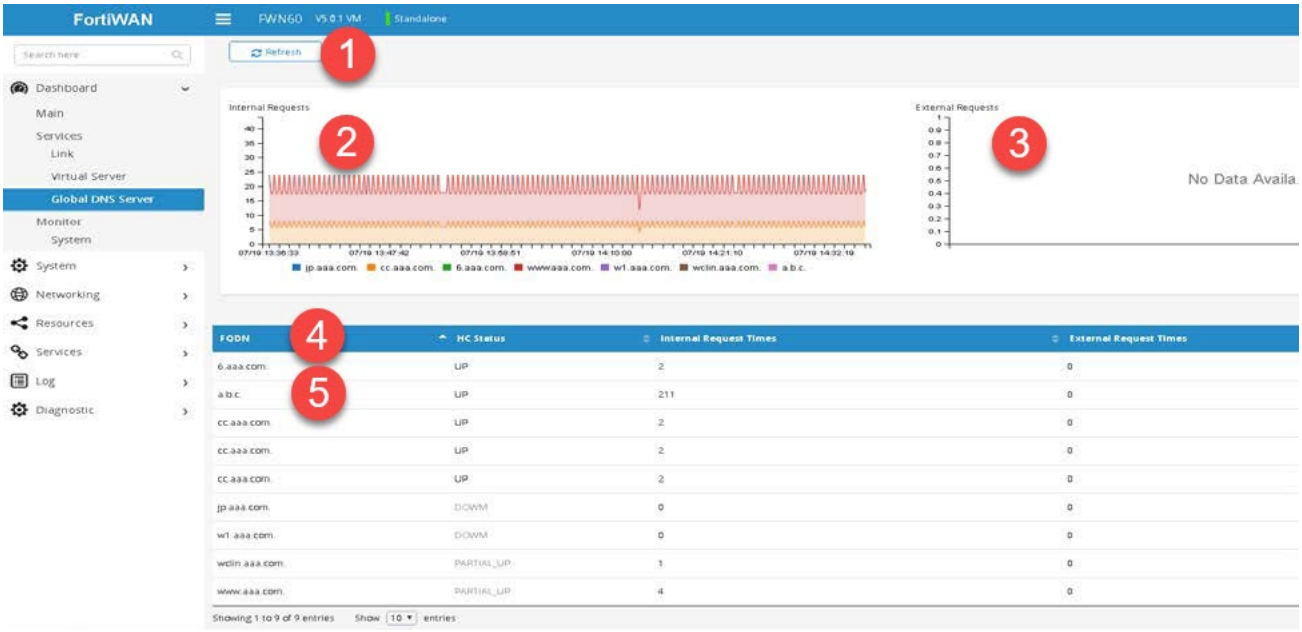

1. Click **Refresh** to update the live data.
2. Shows the internal requests. The legend below the graph corresponds to the port names.
3. Shows the external requests. The legend below the graph corresponds to the port names.
4. Click a column heading to sort in alpha or reverse alpha order.
5. Click a port name to view details. A new page appears in which you can zoom in for a more granular view.

Monitor

The Monitor dashboards allow you to monitor your systems.

DHCP Monitor

This page shows the IP addresses assigned when local FortiWAN is set up as a DHCP server. Server name, IP address, MAC and Expires can be displayed.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Monitor > System**, then select the **DHCP Monitor** tab.

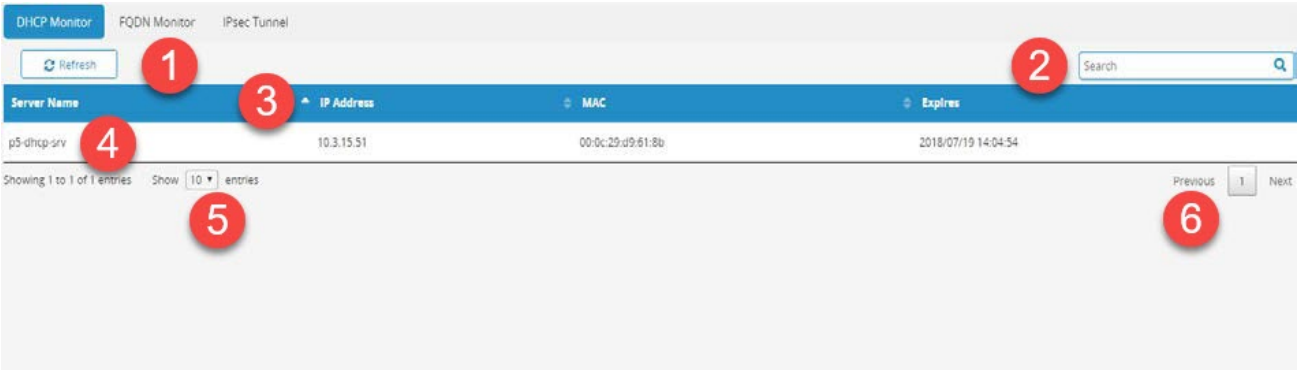

1. Click **Refresh** to update the live data.
2. Search for a port name to filter your view.
3. Click a column heading to sort in alpha or reverse alpha order.
4. Click a port name to view details. A new page appears, in which you can zoom in for a more granular view.
5. Select how many entries to view on a page. You can choose 10, 25, or 50.
6. Click **Previous** or **Next** to see more pages. This only works if you have multiple pages.

FQDN Monitor

This page shows Hit numbers and expire time for FQDN entries configured in **Resources > Address > Address** of FQDN types and searched by local system.

Access

- From the Dashboard, go to **Monitor > System**, then select the **FQDN Monitor** tab.

1. Click **Refresh** to update the live data.
2. Search for a port name to filter your view.
3. Click a column heading to sort in alpha or reverse alpha order.
4. Click a port name to view details. A new page appears, in which you can zoom in for a more granular view.

5. Select how many entries to view on a page. You can choose 10, 25, or 50.
6. Click **Previous** or **Next** to see more pages. This only works if you have multiple pages.

IPSec Tunnel

The IPSec Tunnel page shows the status and incoming and outgoing statistics of IPSec tunnels.

Access

- From the Dashboard, go to **Monitor > System**, then select the **IPSec Tunnel** tab.

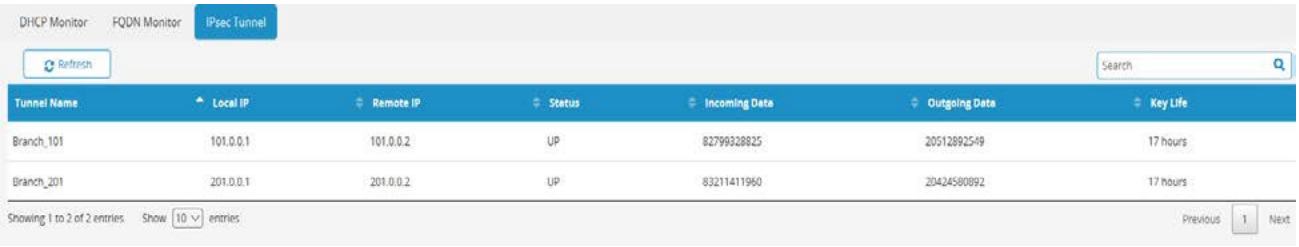

System

These topics describe the System tabs in the FortiWAN interface.

Administrator settings

These topics describe administrator settings from the System menu.

Administrator access

- As soon as possible during initial setup, give the default administrator, admin, a password. This super- administrator account has the highest level of permissions possible, and access to it should be limited to as few people as possible.
- Change all administrator passwords regularly. Set a policy—such as every 60 days—and follow it. (Mark the Change Password check box to reveal the password dialog.)
- Instead of allowing administrative access from any source, restrict it to trusted internal hosts. On those computers that you have designated for management, apply strict patch and security policies. Always password-encrypt any configuration backup that you download to those computers to mitigate the information that attackers can gain from any potential compromise.
- Don't use the default administrator access profile for all new administrators. Create one or more access profiles with limited permissions tailored to the responsibilities of the new administrator accounts.
- By default, an administrator logged in who’s idle for more than 30 minutes times, is logged out. While not recommended, you can change this to a longer period in Timeout. Left unattended, a UI or CLI session could allow anyone with physical access to your computer to change system settings. Small idle timeouts mitigate this risk.
- Administrator passwords should be at least 8 characters long and include both numbers and letters.
- Restrict administrative access to a single network interface (usually port 1), and allow only the management access protocols needed.
- Use only the most secure protocols. Disable ping, except during troubleshooting. Disable HTTP, SNMP, and Telnet unless the network interface only connects to a trusted, private administrative network.
- Disable all network interfaces that should not receive any traffic.
    - For example, if administrative access is typically through port 1, the Internet is connected to port 2, and servers are connected to port 3, you would disable (“bring down”) port 4.

This would prevent an attacker with physical access from connecting a cable to port 4 and thereby gaining access if the configuration inadvertently allows it.

- Immediately revoke certificates that have been compromised. If possible, automate the distribution of certificate revocation lists.

Admin

The **Admin** tab lets you add users, assign access profiles, and assign dashboard views to a user.

Define Access Profiles before defining users.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > Administrator**, then select the **Admin (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
Name	Enter a unique user name. Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Password	Enter a password for this user name.
Trusted Hosts	Enter the IP address of the trusted host.
Profile	Choose from the drop down list. See Access profiles on page 35.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

To change a user password

1. Click the key icon  at the end of the row for the user admin to open the change password editor.
2. Enter the new password, then enter again to confirm, and **Save**.

Administrator users

We recommend that only network administrators—and if possible, only a single person—use the admin account. You can configure accounts that provision different scopes of access. For example, you can create an account for a security auditor who must only be able to view the configuration and logs, but not change them.

Prerequisites

- If you want to use RADIUS or LDAP authentication, you must have already have created the RADIUS server or LDAP server configuration.
- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > Administrator**, then select the **Admin (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Name of the administrator account, such as admin1 or admin@example.com. If you use LDAP or RADIUS, specify the LDAP or RADIUS user name. This is the user name that the administrator must provide when logging in to the CLI or UI. The users are authenticated against the associated LDAP or RADIUS server. Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Password	Set a strong password for all administrator accounts. The password should be at least eight characters long, be sufficiently complex, and be changed regularly. To check the strength of your password, you can use a utility such as Microsoft's password strength meter.
Trusted Hosts	Source IP address and netmask from which the administrator is allowed to log in. For multiple addresses, separate each entry with a space. You can specify up to three trusted areas. They can be single hosts, subnets, or a mixture. Configuring trusted hosts hardens the security of the system. In addition to knowing the password, an administrator must connect only from the computer or subnets you specify. Trusted host definitions apply both to the UI and to the CLI when accessed through Telnet, SSH, or the CLI console widget. Local console access is not affected by trusted hosts, as the local console is by definition not remote, and does not occur through the network. If ping is enabled, the address you specify here is also a source IP address to which the system will respond when it receives a `ping` or `traceroute` signal. To allow logins only from one computer, enter only its IP address and 32- or 128- bit netmask: 192.0.2.1/32 2001:0db8:85a3:::8a2e:0370:7334/128

Setting	Guidelines
	To allow login attempts from any IP address (not recommended), enter: 0.0.0.0/0  If you restrict trusted hosts, do so for all administrator accounts. Failure to do so means that all accounts are still exposed to the risk of brute force login attacks. This is because if you leave even one administrator account unrestricted (i.e. 0.0.0.0/0), the system must allow login attempts on all network interfaces where remote administrative protocols are enabled, and wait until after a login attempt has been received in order to check that user name's trusted hosts list. • To allow login from the Internet, set a longer and more complex New Password, and enable only secure administrative access protocols. We also recommend that you restrict trusted hosts to IPs in your administrator's geographical area. • For improved security, restrict all trusted host addresses to single IP addresses of computer(s) from which only this administrator will log in.
Profile	Select a user-defined or predefined profile. The predefined profile named *super_ admin_prof* is a special access profile used by the admin account. However, selecting this access profile will not confer all permissions of the admin account. For example, the new administrator would not be able to reset lost administrator passwords.  This option does not appear for the **admin** administrator account, which always uses the *super_admin_prof* access profile.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.

- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** ✖ icon at the end of the row.
- To change the password, click the **Key** 🔑 at the end of the row, then enter the new password.

Password policies

A password policy is a set of rules designed to enhance computer security. A good password policy encourages users to create strong passwords and use them properly. For your network and data security and integrity, we strongly recommend the enforcement of strong password policies when using FortiWAN.

You can set a password policy to be enforced for all users.

Access

- From the Dashboard, go to **System > Admininstrator**, then select the **Password Policy** tab.

To set a password policy

1. Enter your settings:

Setting	Guidelines
Password Policy	Toggle ON enable password policy.
Minimum Length	Enter the minimum length of password, from 8 (default) to 32 characters.
Must Contain	Select any options you want to enforce: • Upper Case Letter • Lower Case Letter • Number • Non Alphanumeric (for example, symbols)

2. **Save.**

Access profiles

Access profiles assign permissions to roles.

When a user only has read access to a feature, they can access the web page for that feature, and can use the get and show CLI command for that feature, but can't make changes to the configuration.

Access profiles often reflect the specific job of each user (role), such as account creation or log auditing. Access profiles can limit each user account to their assigned role. This is sometimes called role-based access control (RBAC).

For complete access to all commands and abilities, you must log in with the account name **admin**.

Prerequisites

- Read-Write permission for System settings.

The **super_admin_prof** access profile is a special access profile assigned to and required by the admin account. This default profile and can't be changed or deleted. This profile has permissions similar to the UNIX root account.

Access

- From the Dashboard, go to **System > Administrator**, then select the **Access Profile** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
System Networking Server Load Balance Firewall Log Link Load Balance Global DNS Server Shared Resources	Select the setting allowed for the user. • **Read** - view access, issue CLI command • **Read-Write** - view, change, and execute access, issue CLI command • **No access** - not visible to the user

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row. Doesn't apply to the primary admin user.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row. Doesn't apply to the primary admin user.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Basic

Assign basic settings that apply to all users from this page.

Access

- From the Dashboard, go to **System > Administrator**, then select the **Basic** tab.

Prerequisites

- Read-Write permission for System settings.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
host name	Configure a host name to facilitate system management. If you use SNMP, for example, the SNMP system name is derived from the configured host name. The host name can be up to 35 characters in length. It can include US-ASCII letters, numbers, hyphens, and under- scores, but not spaces and special characters. The System Information widget and the get system status CLI command display the full host name. If the host name is longer than 16 characters, the name is truncated and ends with a tilde (~) to indicate that additional characters exist, but are not displayed.
Default Certificate	Enter the default certificate. By default this is set as Factory.
Idle Timeout	Log out an idle administrator session. The default is 30 minutes.
HTTPS Port	Enter the port for the HTTPS service. Usually, HTTPS uses port 443.
HTTP Port	Enter the port for the HTTP service. Usually, HTTP uses port 80.
SSH Port	Enter the port for the SSH service. Usually, SSH uses port 22.
Telnet Port	Enter the port for the Telnet service. Usually, Telnet uses port 23.
Language	English or Simplified Chinese.
Hardware SSL	Toggle **ON** to enable.
GUI System	Enable the UI system. This can’t be configured.
GUI Router	Enable the UI router. This can’t be configured.
GUI Log	Enable the UI log. This can’t be configured.
SSH CBC Cipher	Toggle **ON** to enable.
SSH HMAC MD5	Toggle **ON** to enable.

Setting	Guidelines
Config Sync	Toggle **ON** to enable. This feature is related to Pushing/Pulling configurations, not HA synchronization. Disabled by default.
Total Sessions	Enter sessions allowed to access to the UI. The default is 10. The valid range is 0 to 128.
Device	Enter the device ID.
Tunnel Quality Check	• **Enable** - turn on VPN tunnel quality detection • **Disable** - turn off VPN tunnel quality detection
Local MNT Traffic Guarantee BM	• **Enable** - turn on local management traffic guarantee (via Bandwidth Management). • **Disable** - turn off local management traffic guarantee (via Bandwidth Management).

Maintenance

Set system time from the Maintenance tab. Scheduling, logging, and SSL/TLS-related features depend on system time.

Use Network Time Protocol (NTP) to maintain the system time. When NTP isn't available or is impractical, you can set the system time manually.

Access

- From the Dashboard, go to **System > Administrator**, then select the **Maintenance** tab.

Prerequisites

- Read-Write permission for System settings.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
System Time	Displays the system time. You can use NTP to set the system time, or use the controls to set the system time manually. Enter time in HH:MM:SS format.
Daylight Saving Time	Enable to have the system adjust its own clock when its time zone changes between daylight saving time (DST) and standard time.
Time Zone	Select the time zone where the appliance is located.
NTP	Toggle ON to enable using an NTP server.
NTP Server	Enter a space-separated list of IP addresses or FQDNs for an NTP server or pool, such as `pool.ntp.org`. To find an NTP server, go to http://www.ntp.org.
Synchronizing Interval	Enter the number of minutes the system should synchronize the time with the NTP server. The default is 60 minutes. The range is 1-1440.

For information on Firmware, see Updating firmware on page 40.

You can also set system time using CLI commands. See CLI: Set System Time on page 236.

Updating firmware

Consider the following to help you determine whether to follow a standard or non-standard upgrade procedure:

- **HA** - Updating firmware on an HA cluster requires additional steps for a standalone appliance. See Updating firmware for an HA cluster on page 256.
- **Re-imaging** - When installing a firmware version that requires a different size system partition, you might need to re-image the boot device. Consult the release notes. In that case, don't install the firmware using this procedure. Instead, see Restoring firmware (“clean install”) on page 278.
- **Downgrades** - If you are downgrading the firmware to a previous version, and the settings are not fully backwards compatible, the system might remove incompatible settings or use the default values for that version of the firmware. You might need to reconfigure some settings.

Important: Read the release notes for release-specific upgrade considerations.

Firmware can be loaded on two disk partitions: the active partition and the alternate partition. The upgrade procedure:

- Updates the firmware on the inactive partition and then makes it the active partition.
- Copies the firmware on the active partition, upgrades it, and installs it in place of the configuration on the inactive partition.

For example, if partition 1 is active, and you perform the upgrade procedure:

- Partition 2 is upgraded and becomes the active partition; partition 1 becomes the alternate partition.
- The configuration on partition 1 remains in place; it is copied, upgraded, and installed in place of the configuration on partition 2.

This preserves the working system state in the event the upgrade fails or is aborted.

Prerequisites

- Download the firmware file from the Fortinet Customer Service & Support website: https://support.fortinet.com/
- Read the release notes for the version you plan to install.
- **Back up your configuration** before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- Superuser permission (user admin).

Access

- From the Dashboard, go to **System > Administrator**, then select the **Maintenance** tab.

Boot the firmware on the alternate partition

- Click **Boot Alternate Firmware**.

The system reboots, the alternate becomes the active firmware, and the active becomes the alternate firmware.

Update the firmware

- Go to the bottom of the page to see the **Upgrade** section.
- (optional) **HA Sync** - Toggle **ON** to use.
- Click **Choose File** to navigate to the file, then click **Open**.
- Click the **arrow** to upload

  —OR—

  Click the **X** to cancel the upload.

  The system replaces the firmware on the alternate partition and reboots. The alternate (upgraded) partition becomes the active, and the active becomes the alternate.

When you update software, you are also updating the UI. To ensure the UI updates correctly:

- Clear your browser cache.
- Refresh the page.

In most environments, press Ctrl-F5 to force the browser to get a new copy of the content from the application. See the Wikipedia article on browser caching issues for a summary of tips for many environments: https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache.

You can also update the firmware in the CLI. See CLI: Update firmware on page 238.

Back up and restore

You can save a copy of the configuration. A simple backup file is a text file. A full backup is a .TAR file, and includes the complete configuration files, plus any files you have imported, including error page files, script files, and ISP address book files. Configuration backups don't include data such as logs and reports.

Backups can be used to:

- Save the configuration as CLI commands that a co-worker or Fortinet support can use to help you resolve issues with misconfiguration.
- Restore the system to a known functional configuration.

- Create a template configuration you can edit and then load into another system using the restore procedure.

In the event that FortiWAN experiences hardware failure, being able to restore the entire backup configuration minimizes the time to reconfigure a replacement.

Back up files can include sensitive information, such as HTTPS certificate private keys. We strongly recommend that you password-encrypt backup files and store them in a secure location.

Prerequisites

- When restoring a configuration, know its management interface configuration in order to access the UI after the restore procedure is completed. Open the configuration file and make note of the IP address and network requirements for the management interface (port1).
- Know the administrator user name and password.
- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > Administrator**, then select the **Backup & Restore** tab.

Settings

- Set your options.

Setting	Guidelines
Mode	Choose an option: • **Back Up** - to create a backup text file. • **Restore** - to start the restore procedure. Your browser uploads the configuration file and the system restarts with the new configuration. Time required to restore varies by the size of the file and the speed of your network connection. Your UI session is terminated when the system restarts. To continue using the UI, refresh the page and log in again. If the restored system has a different management interface configuration than the previous configuration, you must access the UI using the new management interface IP address.
Storage	• **Local PC** - Back up to the local PC. • **WAN** - Back up directly to the FortiWAN device.
Entire Configuration	Check to include error page files, script files, and ISP address book files. This backup is a tar file.
File	Click **Choose File** to navigation to the file to restore. The option applies to restore operations from the local PC only.

- To **Back Up**, click **Back Up** to start the back up.
- To **Restore**, navigate to the file you want to restore, then **upload** .

System Mail Server

Set up a mail server to receive email notifications sent by the system.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > Administrator**, then select the **System Mail Server** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
Address	Enter the email address to receive notifications.
Port	Enter the port number. The default is 25.
Auth	Toggle ON to authorize the mail server.
Security	• Starttls • None
Username	Enter the user name for the email address.
Password	Enter the password for the email address.

High Availability (HA) settings

Deploying high availability (HA) is highly recommended. When setting up a cluster:

- Isolate HA interface connections from your overall network.

Heartbeat and synchronization packets contain sensitive configuration information and can consume considerable network bandwidth. For best results, directly connect the two HA interfaces using a crossover cable. If your system uses switches instead of crossover cables to connect the HA heartbeat interfaces, those interfaces must be reachable by Layer 2 multicast.

- When configuring an HA pair, pay close attention to the options ARP Packet Numbers and ARP Packet Interval.
- The FortiWAN appliance broadcasts ARP packets to the network to ensure timely fail over. Delayed broadcast intervals can slow performance. Set the value of ARP Packet Numbers no higher than needed.
- When the FortiWAN appliance broadcasts ARP packets, it does so at regular intervals. For performance reasons, set the value for ARP Packet Interval no greater than required.

- Some experimentation might be needed to set these options at their optimum value.
- We recommend that you configure an SNMP community and enable the HA heartbeat failed option to generate a message if the HA heartbeat fails.

Currently, FortiWAN only supports HA configurations for IPv4 address mode; HA isn't supported on IPv6.

HA system requirements

- Appliances must have the same hardware model and same firmware version.
- Redundant network topology: if an active node fails, physical network cabling and routes must be able to redirect traffic to the other member nodes.
- At least one physical port on both HA appliances to be used for heartbeat and data traffic between cluster members. For active-passive fail over pairs, you can connect the ports directly over cable.
- Heartbeat and synchronization traffic between cluster nodes occur over the physical network ports that you designate. If switches are used to connect the nodes, the interfaces must be reachable by Layer 2 multicast.
- Each appliance must be licensed.

Prerequisites

- Read-Write permission to items in the System category.

Access

- From the Dashboard, go to **System > High Availability**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- After saving the configuration, cluster members send heartbeat traffic to each other. Members with the same Group ID join the cluster. They send synchronization traffic through their data links.

Setting	Guidelines
Cluster Mode	Select the mode to set.

Setting	Guidelines
	• Active-Passive • Standalone
Group Name	(optional) Type a name to identify the HA cluster if you have more than one. This setting does not affect HA function. The maximum length is 63 characters.
Group ID	Enter a number to identify the HA cluster. Nodes with the same group ID join the cluster. If you have more than one HA cluster on the same network, each cluster must have a different group ID. The group ID is used in the virtual MAC address that's sent in broadcast ARP messages. The valid range is 0 to 31. The default value is 0.
Priority	(optional) Enter a number indicating priority of the member node when electing the cluster primary node. Smaller numbers have higher priority. The default is 5. The valid range is 0 to 9. By default, unless you enable **Override** (below), up time is more important than this setting.
Config Priority	Enter a priority number. The default value is 100. The valid range is 0 to 255. This setting determines the order the system uses each HA node. Lowest configuration priority takes precedence. When the values are identical on two nodes (whether by default or by configuration), the system uses the configuration of the appliance with the larger serial number to override that of the appliance with the smaller serial number.
Monitor Interface	After HA has been configured on all appliances, you can select the ports to be monitored. One or more network interfaces that correlate with a physical link. These ports will be monitored for link failure. Port monitoring (also called interface monitoring) monitors physical network ports to verify that they are functioning properly and

Setting	Guidelines
	linked to their net- works. You can monitor physical interfaces and 802.3ad aggregated interfaces.  To prevent an unintentional fail over, don't configure port monitoring until you configure HA on all appliances and have plugged in the cables to link the physical network ports to monitor.
Heartbeat Interface	Set the network interface to be used for heartbeat packets. You can configure one or two heartbeat ports. Use the same port number for all cluster members. For example, if you select port3 on the primary node, select port3 as the heartbeat interface on the other member nodes. If a switch is used to connect the heartbeat interfaces, the heartbeat interfaces must be reachable by Layer 2 multicast
Override	Toggle **ON** to make Device Priority a more important factor than up time when selecting the primary node.
Heartbeat Interval	Number of 100-millisecond intervals at which heartbeat packets are sent. This is also the interval at which a node expects to receive heartbeat packets. This part of the configuration is pushed from the primary node to member nodes. The default is 2. The valid range is 1 to 20 (that's, between 100 and 2,000 milliseconds). Although this setting is pushed from the primary node to member nodes, you should initially configure all nodes with the same Detection Interval to prevent inadvertent fail over from occurring before the initial synchronization.
Lost Heartbeat Threshold	Number of times a node retries the heartbeat and waits to receive HA heartbeat packets from the other nodes before concluding the other node is down. This part of the configuration is pushed from the primary node to member nodes. Normally, you don't need to change this setting. Exceptions include: • Increase the failure detection threshold if a failure is detected

Setting	Guidelines
	when none has actually occurred. For example, in an active-passive deployment, if the primary node is very busy during peak traffic times, it might not respond to heartbeat packets in time, and a standby node might assume that the primary node has failed. • Decrease the failure detection threshold or detection interval if administrators and Threshold HTTP clients have to wait too long before being able to connect through the primary node, resulting in noticeable down time. The valid range is from 1 to 60. Although this setting is pushed from the primary node to member nodes, you should initially configure all nodes with the same HB Lost Threshold to prevent inadvertent fail over from occurring before the initial synchronization.
ARP Times	Enter the number of times that the cluster member broadcasts extra address resolution protocol (ARP) packets when it takes on the primary role. Normally, you don't need to change this setting. Exceptions include: • Increase the number of times the primary node sends gratuitous ARP packets if an active-passive cluster takes a long time to fail over or to train the network. Sending more gratuitous ARP packets may help the fail over to happen faster. • Decrease the number of times the primary node sends gratuitous ARP packets if the cluster has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them might generate a large amount of network traffic. As long as the active-passive cluster fails over successfully, you can reduce the number of times gratuitous ARP packets are sent to reduce the amount of traffic produced by a fail over. The valid range is 1 to 60. The default is 5. Although a new NIC has not actually been connected to the network, the member does this to notify the network that a new physical port has become associated with the IP address and

Setting	Guidelines
	virtual MAC of the HA cluster. This is also called “using gratuitous ARP packets to train the network,” and can occur when the primary node is starting up, or during a fail over. Also configure ARP Packet Interval.
ARP Interval	Enter the number of seconds to wait between each broadcast of ARP packets. Normally, you don't need to change this setting. Exceptions include: • Decrease the interval if an active-passive cluster takes a long time to fail over or to train the network. Sending ARP packets more frequently may help the fail over to happen faster. • Increase the interval if the cluster has a large number of VLAN interfaces and virtual domains. Because gratuitous ARP packets are broadcast, sending them might generate a large amount of network traffic. As long as the active-passive cluster fails over successfully, you can increase the interval between when gratuitous ARP packets are sent to reduce the rate of traffic produced by a fail over. The valid range is from 1 to 20. The default is 6 seconds.
Remote IP Monitor	Toggle **ON** to actively monitor remote beacon IP addresses to determine network path availability.
Fail over Threshold	Enter the number of consecutive times that the remote IP address is unreachable to indicate failure. The default is 5. The valid range is 1-300.
Fail over Hold Time	If fail over occurs due to a remote IP monitor test, and this node's role changes (to master or child), it can't change again until the hold time elapses. Hold time can be used to prevent looping. The default hold time is 120 seconds. The valid range is 60-86400.
Auto Config Sync	Toggle **ON** to enable automatic configuration synchronization. When enabled, synchronization occurs immediately when an appliance joins the cluster, and thereafter every 30 seconds. Disable if you want to manually manage synchronization.

Setting	Guidelines
Remote IP Monitor List	
Name	Configuration name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Reference this name in the virtual server configuration. After initially saving the configuration, you can't edit the name.
Remote Address	Remote address to ping.
Source Port	Interface to send the health check ping.
Health Check Interval	Seconds between each health check. Should be more than the timeout to prevent overlapping health checks. The default is 10.
Health Check Timeout	Seconds to wait for a reply before assuming that the health check has failed. The default is 5.
Health Check Retry	Number of retries to confirm up or down. The default is 3 retries. The valid range is 1-10.

FortiGuard

FortiGuard periodically updates the Geo IP Database. Go to the FortiGuard website (https://fortiguard.com) to download the update packages that you can upload to FortiWAN, or you can schedule automatic updates.

Prerequisites

- To perform a manual update, you must download the update file from the FortiGuard website.
- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > FortiGuard**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
FortiCare Support	
FortiCare Support	Contains license information. To change registration or if your license is about to expire, contact Fortinet Service & Support. If your license is invalid, FortiGuard does not send updates to the FortiWAN. The functionality on FortiWAN remains intact and useful, but it is out- of-date.
IP Reputation DB	The FortiGuard IP Reputation service provides a database of known compromised or malicious client IP addresses. The database is updated periodically.
Geo IP Database	Review the version information. To update, see Main on page 23.
Internet Service DB	Supports three Destination Address Types - IP range, IP netmask, and FQDN.
SD-WAN	
Firewall	Appears when the FDS (FortiGuard Distribution Servers) contract includes a firewall license.
Update Schedule	
Scheduled Update	Toggle **ON** to enable updates.
Scheduled Update Frequency	Select a frequency. • **Every** - Schedule periodic updates. Enter the time interval to perform updates. • **Daily** - Schedule daily updates. Enter the time of day to perform the update. • **Weekly** - Schedule weekly updates. Enter the day and time to perform the update.
Scheduled Update Day	Choose a day from the drop down list.
Scheduled Update Time	Enter a time to update. For example, HH:MM.HH is 0-23. MM is 0, 15, 30, or 45.
Override Server	If you are unable to make connections to the standard FortiGuard

Setting	Guidelines
	server, toggle **ON** to enable connection to the override server address given to you by Fortinet Service & Support.
Override Server Address	Enter the Override server IP address.

SNMP settings

Many organizations use SNMP (simple network management protocol) to track the health of their systems. FortiWAN supports SNMP v1, v2c, and v3.

SNMP depends on network devices that maintain standard management information bases (MIBs). MIBs describe the structure of the management data maintained on the device. Some MIB definitions are standard for all network devices, and some are vendor and product-family specific.

The FortiWAN system runs an SNMP agent to communicate with the SNMP manager. The agent enables the system to respond to SNMP queries for system information and to send SNMP traps (alarms or event messages) to the SNMP manager.

Figure 8 - SNMP communication

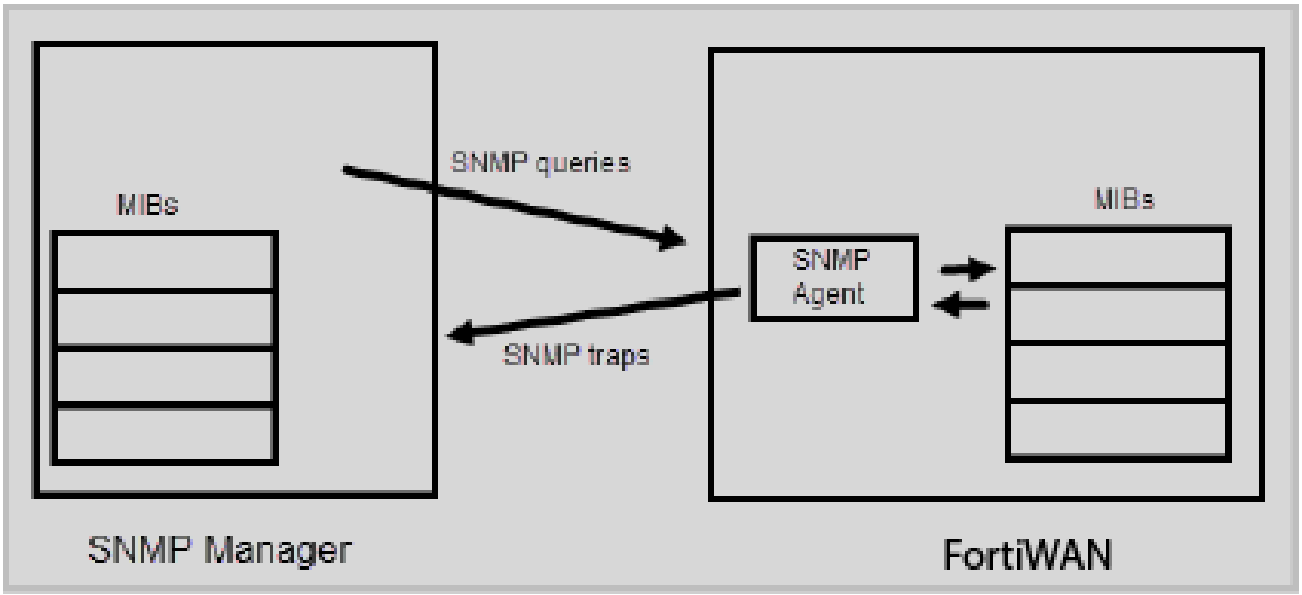

With SNMP v1 and v2c managers, configure SNMP communities to connect FortiWAN and the SNMP manager. The SNMP Manager sends the community string along with all SNMP requests. If the community string is correct, the device responds with the requested information. If the community string is incorrect, the device simply discards the request and does not respond.

Fortinet strongly recommends that you don't add FortiWAN to the community named public. This default name is well-known, and attackers that attempt to gain access to your network often try this name first.

With SNMPv3 managers, configure SNMP users to connect FortiWAN and the SNMP manager. Queries and traps include username/password authentication, along with an encryption key. FortiWAN implements the user security model described in RFC 3414.

Prerequisites

- On the SNMP manager, you must verify that the SNMP manager is a member of the community to which the FortiWAN system belongs, and you must compile the necessary Fortinet-proprietary management information blocks (MIBs) and Fortinet-supported standard MIBs. For more information, see Fortinet MIBs on page 282.
- In the FortiWAN interface settings, you must enable SNMP access on the network interface through which the SNMP manager connects.
- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > SNMP**, then select the **System Information (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
SNMP Agent	Toggle **ON** to activate the SNMP agent, so the system can send traps and receive queries.
Description	A description or comment about the system, such as “don’t reboot”. The description can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Contact	Contact information for the administrator or other person responsible for this system, such as a phone number (555-5555) or name (jdoe). The contact information can be up to 35 characters long, and can

Setting	Guidelines
	contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).
Location	Physical location of the appliance, such as “floor_2”. The location can be up to 35 characters long, and can contain only letters (a-z, A-Z), numbers, hyphens (-) and underscores (_).

See Download SNMP MIBs on page 55 for more information on the bottom section of the page.

Download SNMP MIBs

FortiWAN allows you to download full FortiWAN and Fortinet Core MIB files, which provides more options for server load balance, global serer load balance, link load balance, and firewall with SNMP traps.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > SNMP**, then select the **System Information (default)** tab.

To download an SNMP MIB file

- At the bottom of the page, click **Download FortiWAN MIB File** or **Download Fortinet Core MIB File**.

  The file automatically downloads to your local drive.

SNMP threshold

Access

- From the Dashboard, go to **System > SNMP**, then select the **Threshold** tab.

Prerequisites

- Read-Write permission for System settings.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
CPU	• **Trigger** - The default is 80% utilization. • **Threshold** - The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period. • **Sample Period** - The default is 600 seconds. The default is 30 seconds.
Memory	• **Trigger** - The default is 80% utilization. • **Threshold** - The default is 3, meaning the event is reported when the condition has been triggered 3 times in a short period. • **Sample Period** - The default is 600 seconds. The default is 30 seconds.
Logdisk	• **Trigger** - The default is 90% utilization. • **Threshold** - The default is 1, meaning the event is reported each time the condition is triggered. • **Sample Period** - The default is 7200 seconds. The default is 3600 seconds.

Configure SNMP v1/v2

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > SNMP**, then select the **SNMPv1/v2** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
SNMP Community Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. Name of the SNMP community to which the FortiWAN system and at least one SNMP manager belongs, such as management. You must configure the FortiWAN system to belong to at least one SNMP community so that community's SNMP managers can query system information and receive SNMP traps. You can add up to three SNMP communities. Each community can have a different con- figuration for queries and traps, and the set of events that trigger a trap. You can also add the IP addresses of up to eight SNMP managers to each community to des- ignate the destination of traps and which IP addresses are permitted to query the FortiWAN sys- tem.
Status	Toggle **ON** to enable the configuration.
SNMP v1 Status	Toggle **ON** to enable the SNMP v1 configuration.
SNMP v1 Port	Enter the port number on which the system listens for SNMP v1 queries from the SNMP managers in this community. The default is 161. The range is 0 to 65535.
SNMP v2 Status	Toggle **ON** to enable the SNMP v2 configuration.
SNMP v2 Port	Enter the port number on which the system listens for SNMP v2 queries from the SNMP managers in this community. The default is 161. The range is 0 to 65535.
Trap v1 Status	Toggle **ON** to enable the Trap v1 configuration.
Trap v1 Local Port	Enter the SNMP Trap v1 local port number. The default is 162. The range is 0 to 65535.
Trap v1 Remote Port	Enter the SNMP Trap v1 remote port number. The default is 162. The range is 0 to 65535.

Setting	Guidelines
Trap v2 Status	Toggle **ON** to enable the Trap v2 configuration.
Trap v2 Local Port	Enter the SNMP Trap v2 local port number. The default is 162. . The range is 0 to 65535.
Trap v2 Remote Port	Enter the SNMP Trap v2 remote port number. The default is 162. . The range is 0 to 65535.
Events	Select to enable SNMP event reporting for the following thresholds: • **CPU** - CPU usage has exceeded 80%. • **Memory** - Memory (RAM) usage has exceeded 80%. • **Log disk usage** - Disk space usage for the log partition or disk has exceeded 90%. • **Platform** - Reserved for future use. • **Security** - Reserved for future use. • **Application** - Reserved for future use.
Host	Appears after the configuration is initially saved. Click **Add** to open the configuration editor. • **IP Address** - Subnet address for the SNMP manager to receive traps and be permitted to query the FortiWAN system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiWAN system, enter 0.0.0.0/0. For security best practice reasons, however, this isn't recommended. • **Host Type** - Whether the host can send queries, receive traps, or any (both).  The system sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment. If there are no other host IP entries, entering only 0.0.0.0/0

Setting	Guidelines
	effectively disables traps because there is no specific destination for trap packets. If you don't want to disable traps, you must add at least one other entry that specifies the IP address of an SNMP manager. Test both traps and queries (assuming you have enabled both). Traps and queries typ- ically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional. To test queries, from your SNMP manager, query the FortiWAN appliance. To test traps, cause one of the events that should trigger a trap.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Configure SNMP v3

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **System > SNMP**, then select the **SNMPv3** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. This user name is what the SNMP Manager uses to communicate with the SNMP Agent.
Status	Toggle **ON** to enable the configuration.
Security Level	• **No Auth And No Privacy** - don't require authentication or encryption. • **Auth But No Privacy** - Authentication based on MD5 or SHA algorithms. Select an algorithm and specify a password. • **Auth And Privacy** - Authentication based on MD5 or SHA algorithms, and encryption based on AES or DES algorithms. Select an Auth Algorithm and specify an Auth Password; and select a Private Algorithm and specify a Private Password.
Queries Status	Toggle **ON** to enable queries for SNMP v3.
SNMP v3 Port	Enter the port number on which the system listens for SNMP queries from the SNMP managers for this user. The default is 161.
Trap Status	Toggle **ON** to enable traps for SNMP v3.
Trap Local Port	Enter the source (Local) port number for trap packets sent to SNMP managers for this user. The default is 162.
Trap Remote Port	Enter the destination (Remote) port number for trap packets sent to SNMP managers for this user. The default is 162.
Events	Choose SNMP event reports for the following thresholds: • **CPU** - CPU usage has exceeded 80%.

Setting	Guidelines
	• **Memory** - Memory (RAM) usage has exceeded 80%. • **Log disk usage** - Disk space usage for the log partition or disk has exceeded 90%. • **Platform** - Reserved for future use. • **Security** - Reserved for future use. • **Application** - Reserved for future use.
Host	Appears after the configuration is initially saved. Click **Add** to open the configuration editor. • **IP Address** - Subnet address for the SNMP manager to receive traps and be permitted to query the FortiWAN system. SNMP managers have read-only access. You can add up to 8 SNMP managers to each community. To allow any IP address using this SNMP community name to query the FortiWAN system, enter 0.0.0.0/0. For security best practice reasons, however, this isn't recommended. • **Host Type** - Whether the host can send queries, receive traps, or any (both).  The system sends security-sensitive traps, which should be sent only over a trusted network, and only to administrative equipment. If there are no other host IP entries, entering only 0.0.0.0/0 effectively disables traps because there is no specific destination for trap packets. If you don't want to disable traps, you must add at least one other entry that specifies the IP address of an SNMP manager. Test both traps and queries (assuming you have enabled both). Traps and queries typ- ically occur on different port numbers, and therefore verifying one does not necessarily verify that the other is also functional. To test queries, from your SNMP manager, query the FortiWAN appliance. To test traps, cause one of the events that should trigger a trap.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Networking

Configure basic settings for your network from this part of the UI.

Interface settings

You can edit the physical interface configuration. You can't create or delete a physical interface configuration.

Prerequisites

- Read-write permission for networking settings.

Access

- From the Dashboard, go to **Networking > Interface**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Interface name is defined as a combined string composed of two parts. Example: Test123. "Test" is part1, and "123" is part2. part1 can input character range [a-zA-Z0-9]. part2 can input character range [a-zA-Z0-9._-].
Mode	• **Static** - Enter a static IP address. The IP address must be on the same subnet as the network to which the interface connects. Two network interfaces can't have IP addresses on the same subnet (i.e. overlapping subnets). • **DHCP** - If you configure an interface to use DHCP, FortiWAN automatically broadcasts a DHCP request from the interface. The interface is configured with the IP address, any DNS server addresses, and the default gateway address that the DHCP server provides. • **PPPoE -** Use PPPoE to retrieve a configuration for the IP address, gateway, and DNS server. For example, if this interface uses a DSL connection to the Internet, your ISP may require this option.
Mode - Static	
IPv4/Netmask	Enter the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 192.0.2.5/24. Dotted quad formatted subnet masks are not accepted.
IPv6/Netmask	Enter the IP address and CIDR-formatted subnet mask, separated by a forward slash (/), such as 2001:0db8:85a3::8a2e:0370:7334/64. Dotted quad formatted subnet masks are not accepted.
Relay Enable	Enable / Disable DHCP relay

Setting	Guidelines
Secondary IP	Secondary IP addresses can be used when you deploy the system so that it belongs to multiple logical subnets. When multiple IP addresses are assigned to an interface, you must also assign static addresses to them. To add secondary IP addresses, enable the feature and **Save**. After you have saved it the first time, you can edit it to add secondary IP addresses and enable inbound traffic to that address.
Mode - DHCP	
DNS Server Override	Toggle **ON** to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page.
Mode - PPPoE	
User Name	PPPoE account user name
Password	PPPoE account password.
Discovery Retry Timeout	Seconds the system waits before it retries to discover the PPPoE server. The default is 5 seconds. The valid range is 1-255.
PPPoE Service	Defines the service that the router can provide to a PPPoE client.
PPPoE IP	FortiWAN will ask the PPPoE server for this specific IP.
DNS Server Override	Enable this to use the DNS addresses retrieved from the DHCP server instead of the DNS server settings in the **Networking > DNS** page. (See DNS - Networking on page 79.)
Type	• Physical • VLAN • Aggregate • Loopback • Softswitch

Setting	Guidelines
Type - VLAN	
VLAN ID	VLAN ID of packets that belong to this VLAN. If one physical network port (that's, a VLAN trunk) will handle multiple VLANs, create multiple VLAN subinterfaces on that port, one for each VLAN ID that will be received. The valid range is between 1 and 4094. The value you specify must match the VLAN ID added by the IEEE 802.1q-compliant router or switch connected to the VLAN subinterface.
Interface	Physical interface associated with the VLAN; for example, port2.
Type - Aggregate	
Member	Select the physical interfaces that are included in the aggregation.
Aggregate Mode	Link aggregation type: • 802.3ad • Balance-alb • Balance-rr • Balance-tlb • Balance-xor • Broadcast
Type - Softwitch	
Member	Select the physical interfaces that are included in the aggregation.
General	
Allow Access	Allow inbound service traffic. Select from the following options: • **HTTP** - Enables connections to the UI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer. • **HTTPS** - Enables secure connections to the UI. We recommend this option instead of HTTP.

Setting	Guidelines
	• **Ping** - Enables `ping` and `traceroute` to be received on this network interface. When it receives an ECHO_ REQUEST (“ping”), FortiWAN will reply with ICMP type 0 (ECHO_ RESPONSE or “pong”). • **SNMP** - Enables SNMP queries to this network interface. • **SSH** - Enables SSH connections to the CLI. We recommend this option instead of Telnet. • **Telnet** - Enables Telnet connections to the CLI. We recommend this option only for network interfaces connected to a trusted private network, or directly to your management computer.
MTU	Maximum transmission unit.
Status	The Status column isn't the detected physical link status; it is the administrative status (Up/Down) that indicates whether you permit the network interface to receive and/or transmit packets.
Speed	Auto / 10half / 10full / 100half / 100full / 1000full
Dedicate to Management	Enable / Disable
Zone	Zones are a group of one or more FortiWAN interfaces, both physical and virtual, that you can apply security policies to control inbound and outbound traffic. FortiWAN supports zone types: WAN / LAN / DMZ

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

You can also perform this action using CLI commands. See CLI Configure Network Interfaces on page 237.

Example: How to configure a VLAN interface

1. From the Dashboard, go to **Networking > Interface**, the click **Add**.
2. Enter your options as needed.
   - For **Type**, choose **VLAN**.
   - Enter a **VLAN ID**.
3. **Save**.

You can also perform this action using CLI commands. See CLI: Configure a VLAN interface on page 235.

Types of Interfaces

Physical interfaces

Each physical network port (or, on FortiWAN-VM, a vNIC) has a network interface that directly corresponds to it— that is, a “physical network interface.”

Physical ports have three uses:

- **Management** - The network interface named port1 is typically used as the management interface.
- **HA** - If you plan to deploy HA, you must reserve a physical port for HA heartbeat and synchronization traffic. Do not configure the network interface that is used for HA; instead, leave it unconfigured or “reserved” for HA.
- **Traffic** - The remaining physical ports can be used for your target traffic—these are your “traffic interfaces.”

Traffic interfaces can be associated with logical interfaces. The system supports three types of logical interfaces: VLAN, Aggregate and Software switch.

VLAN interfaces

Use IEEE 802.1q VLAN to reduce the size of a broadcast domain, thereby reducing the amount of broadcast traffic received by network hosts, improving network performance.

Unlike physical LANs, VLANs don't require you to install separate hardware switches and routers to achieve this effect. Instead, VLAN-compliant switches restrict broadcast traffic based upon whether its VLAN ID matches that of the destination network. As such, VLAN trunks can be used to join physically distant broadcast domains as if they were close.

The VLAN ID is part of the tag that's inserted into each Ethernet frame in order to identify traffic for a specific VLAN. FortiWAN appliances handle VLAN header addition automatically, so you don't need to adjust the maximum transmission unit (MTU). Depending on whether the device receiving a packet operates at Layer 2 or Layer 3 of the network, a VLAN tag might be added, removed, or rewritten before forwarding to other nodes on the network. For example, a Layer 2 switch typically adds or removes a tag when forwarding traffic among members of the VLAN, but does not route tagged traffic to a different VLAN ID. In contrast, a FortiWAN content- based routing policy might forward traffic between different VLAN IDs (also known as inter-VLAN routing).

Cisco Discovery Protocol (CDP) is supported for VLANs.

VLANs are not designed to be a security measure, and should not be used where untrusted devices or individuals outside of your organization have access to the equipment. VLAN tags are not authenticated, and can be ignored or modified by attackers. VLAN tags rely on the voluntary compliance of the receiving host or switch.

Aggregate interfaces

Link aggregation (also called NIC teaming-bonding or link bundling) forms a network interface that queues and transmits over multiple wires (also called a port channel), instead of only a single wire (as FortiWAN would normally do with a single network interface per physical port). This multiplies the bandwidth that's available to the network interface, and therefore is useful if FortiWAN is deployed inline with your network backbone.

Link aggregation on FortiWAN complies with IEEE 802.3ad and distributes Ethernet frames using a modified round-robin behavior. If a port in the aggregation fails, traffic is redistributed automatically to the remaining ports with the only noticeable effect being a reduced bandwidth. When broadcast or multicast traffic is received on a port in the aggregation, reverse traffic will return on the same port.

When link aggregation uses a round-robin that considers only Layer 2, Ethernet frames that belong to an HTTP request can sometimes arrive out of order. Because network protocols at higher layers often don't gracefully handle this (especially TCP, which may decrease network performance by requesting retransmission when the expected segment does not arrive), the FortiWAN frame distribution algorithm is configurable. For example, if you notice that performance with link aggregation isn't as high as you expect, you could try configuring FortiWAN to queue related frames consistently to the same port by considering the IP session (Layer 3) and TCP connection (Layer 4), not simply the MAC address (Layer 2).

Also, configure the router, switch, or other link aggregation control protocol (LACP)-compatible device to which FortiWAN is connected with the same speed/duplex settings, and with ports that

can be aggregated. In a deployment like this, the two devices use the cables between the ports to form a trunk, not an accidental Layer 2 (link) network loop. FortiWAN uses LACP to detect the following conditions:

- Suitable links between itself and the other device, and form a single logical link.
- Individual port failure so that the aggregate can redistribute queuing to avoid a failed port.

Loopback interfaces

A loopback interface is a logical interface that's always up (no physical link dependency) and the attached subnet is always present in the routing table.

The IP address of the FortiWAN loopback interface does not depend on one specific external port, and therefore you can access it through several physical or VLAN interfaces.

Loopback interfaces still require appropriate firewall policies to allow traffic to and from the interfaces..

Loopback interfaces are a good practice for OSPF. To make troubleshooting OSPF easier, set the OSPF router ID the same as the loopback IP address, and remember the management IP addresses. You can enable dynamic routing protocols on loopback interfaces.

Software switch interfaces

A software switch, is a virtual switch that's implemented at the software, or firmware level, rather than the hardware level. A software switch can be used to simplify communication between devices connected to different FortiWAN interfaces. For example, using a software switch, you can place the FortiWAN interface connected to an internal network on the same subnet as your wireless interfaces. Then devices on the internal network can communicate with devices on the wireless network without any additional configuration such as additional security policies, on the FortiWAN unit.

This is also useful if you require more hardware ports for the switch on a FortiWAN unit. For example, if your FortiWAN unit has a 4-port switch, WAN1, WAN2 and DMZ interfaces, and you need one more port, you can create a soft switch that can include the 4-port switch and the DMZ interface all on the same subnet. These types of applications also apply to wireless interfaces and virtual wireless interfaces and physical interfaces such as those with FortiWiFi and FortiAP unit.

Similar to a hardware switch, a software switch functions like a single interface. A software switch has one IP address; all of the interfaces in the software switch are on the same subnet. Traffic between devices connected to each interface are not regulated by security policies, and traffic passing in and out of the switch are affected by the same policy.

Consider the following when setting up a software switch:

- Ensure you create a backup of the configuration.
- Ensure you have at least one port or connection such as the console port to connect to the FortiWAN. If you accidentally combine too many ports, you will need a way to undo any errors.
- The ports that you include must not have any link or relation to any other aspect of the FortiWAN unit. For example, DHCP servers, security policies, and so on.
- For increased security, you can create a captive portal for the switch, allowing only specific user groups access to the resources connected to the switch.
- To add an interface to a software switch, the interface can't be referenced by the existing configuration. It must also have its IP address set to `0.0.0.0/0.0.0.0`.

IPsec tunnels

Use IPsec VPN to tunnel traffic between pairs of FortiWAN appliances. See LLB overlay links on page 114.

The overlay link group configuration sets the list of overlay tunnel members, as well as load balancing options, such as algorithm and weight.

When you add IPsec Tunnel configuration, you specify a local and remote IP address. These addresses are IP addresses assigned to a network interface on the local and remote FortiWAN appliance.

If the peer IP is a dynamic IP address (such as DHCP or PPPoE), you must specify Type as equal to dynamic on the local site.

Prerequisites

- Read-Write permission for Networking module settings.

Access

- From the Dashboard, go to **Networking > IPsec VPN**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Local IP	Enter an IP address for the local end of the VPN tunnel. Select one of the following: • **Local IP** - When you select the corresponding interface, the primary IP address of the specified interface will be used. • **Enter** - Enter a IP address of the interface for the Primary IP address or secondary IP address.
Remote IP	The remote VPN gateway's address. Only support IP address.
Key Life	Enter a number that sets a limit on the length of time that a Phase 2 key can be used. The default is 86400 seconds. The range is 120 to 172800 seconds.
DPD	Enable to re-establish VPN tunnels on idle connections and clean up dead IKE peers, as needed.
Proposal Encryp	Choose from the drop down list. • des • 3des • aes128 (default) • aes192 • aes256
Proposal Auth	Choose from the drop down list. • md5 (default) • sha1 • sha256 • sha384 • sha512
DH Group	• **1** - More Modular Exponential (MODP) DH Group with a

Setting	Guidelines
	768-bit modulus. • **2** - MODP with a 1024-bit modulus. • **5** - MODP with a 1536-bit modulus. • **14** - MODP with a 2048-bit modulus.
Psksecret	PreShareKey (PSK). This must be the same at both ends. It encrypts Phase 1 authentication information.
IKE version	• IKE v1 version • IKE v2 version
Type	• **Static** - Peer site use a static IP address. • **Dynamic** - Peer site use a. dynamic IP address, such as DHCP/PPPoE
Peer Type	• **Any** - By default, don't set peer id. • **One** - Set to one, you need to set the peer id.
Local ID	Local identity ID. This Local ID value must match the peer ID value given for th remote VPN peer's Peer Options. The length is from 1 to 254 characters
Interface	The network interface that connects to the other VPN gateway.
Peer ID	Peer identity ID. This Peer ID value must match the Local ID value given for th remote VPN peer's Local ID Options. The length is from 1 to 254 characters

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

DHCP Server settings

Configure a DHCP server from this page.

A DHCP server provides an address, from a defined address range, to a client on the network that requests it. An interface can't provide both a server and a relay for connections of the same type

(regular or IPsec). However, you can configure a regular DHCP server on an interface only if the interface is a physical interface with a static IP address.

You can configure one or more DHCP servers on any FortiWAN interface. A DHCP server dynamically assigns IP addresses to hosts on the network connected to the interface. The host computers must be configured to obtain their IP addresses using DHCP.

If an interface is connected to multiple networks through routers, you can add a DHCP server for each network. The IP range of each DHCP server must match the network address range. The routers must be configured for DHCP relay.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Networking > DHCP Server**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	
Name Guidelines	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Interface	Select the interface to enable DHCP Server on it.
DNS Server	The DNS that FortiWAN responds to the DHCP clients within the DHCP OFFER messages if the clients are sat to automatically get DNS information through DHCP. • **ALL**: answer the DHCP clients with all the defined DNS servers information. • **None**: answer the DHCP clients without containing any DNS server information. • **Enter**: the DNS servers defined in Networking > DHCP Server > DNS Server Address.
DNS Suffix	The domain name suffix that FortiWAN responds to the DHCP clients within the DHCP OFFER messages if the clients are set to automatically get DNS information from DHCP. • **ALL**: answer the DHCP clients with all the defined domain name suffixes. • **None**: answer the DHCP clients without containing any domain name suffixes.
TFTP Server	This option is used to deliver a TFTP server IP address to DHCP clients. When the DHCP server see the request in a DHCP discover from a DHCP client, it returns the TFTP server IP address in its DHCP offer to the client as DHCP option 66. Usually, option 66 is used for IP phone auto- provisioning. You will need to refer to a vender's documentation to configure this option. Enter the IP address or the host name of a TFTP server directly here according to what the device vender provides. FortiWAN DHCP will directly return what is specified here to requests without any encoding/decoding. The DHCP server will ignore the request for option 66 from a DHCP client if this field is leaved

Setting	
	blank.
VSI	DHCP Option 43, Vendor Specific Information.
IP-Range	The address pools that DHCP server assigns and manages IP addresses from. Define the IP ranges by specifying IPv4 Starting Address and IPv4 Ending Address.
Static-Mapping	DHCP server assigns and manages IP addresses according to clients' MAC addresses. An IP address that's mapped to a MAC address is only available to the client with the MAC address. It will not be assigned to other client even it is idle. Define the mapping by specifying MAC Address and the correspondent IPv4 Address.
Client-Mapping	DHCP server assigns and manages IP addresses according to the client ID of DHCP client (the Client Identifier, options code 61, in the options field of DHCP request). An IP address that's mapped to a client ID here is only available to this client. It will not be assigned to other clients even it is idle. Define the mapping by specifying Client ID and the correspondent IPv4 Address. Corresponding setting of client ID on a DHCP client is required.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Links - underlay or overlay

You can use the default link settings, or create your own. You can specify health checks, bandwidth rate thresholds, and spillover threshold behavior for the gateway links you add to link groups.

Prerequisites

- Know the IP addresses of the ISP gateway links used in the network segment where the FortiWAN appliance is deployed.

- Add health check configuration objects that you want to use to check the gateway links. See Links - underlay or overlay on page 76.
- Read-Write permission for Networking module settings.

Access

- From the Dashboard, go to **Networking > Link**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Address	Enter the IP address. Appears when the WAN Type **Static Gateway** is selected. For example, 192.0.2.1
Interface	Select an interface from the drop down list. Appears when the WAN Type **Dynamic Gateway** or **Tunnel** are selected. The default is port1.
WAN Type	•Static Gateway •Dynamic Gateway •Tunnel
Health Check	Toggle **ON** to enable health check. For port1, health check is called LLB_HLTHCK_HOPS LLB_HLTHCK_ICMP For port2, health check is called LLB_HLTHCK_HOPS LLB_HLTHCK_ICMP

Health Check List	Move items to monitor into the Selected Items box. Appears when Health Check is **ON**.
	Health check list supports adding multiple health check objects for one link at the same time. The up/down status depends on the relationship of these health checks:
	*OR
	*AND
	The default health check list for underlay links port1 and port2 include two health checks: LLB_HLTHCK_ICMP, LLB_HLTHCK_HOPS. Their relationship is OR.
	`set health-check-list LLB_HLTHCK_ICMP LLB_HLTHCK_HOPS`
	`set health-check-relation OR`
	That means if any of the health checks succeed, the health

Setting	Guidelines
	check status of the link will be up.
Inbound Bandwidth	Enter the maximum allowed inbound bandwidth in kbps. Maximum bandwidth rate for outbound traffic to this gateway link. If traffic exceeds this threshold, the FortiWAN system considers the gateway to be full and does not dispatch new connections to it. The default is 2000000 Kbps. The valid range is 1 to 2147483647 Kbps.
Outbound Bandwidth	Enter the maximum allowed outbound bandwidth in kbps. The default is 2000000 Kbps. The valid range is 1 to 2147483647 Kbps.
Inbound Spillover Threshold	Enter the maximum inbound bandwidth rate for a link in a spillover load balancing pool. The default is 2000000 Kbps. The valid range is 1 to 2147483647 Kbps.
Outbound Spillover Threshold	Enter the maximum allowed outbound spillover threshold in kbps. The default is 2000000 Kbps. The valid range is 1 to 2147483647 Kbps.
Total Spillover Threshold	Enter the maximum allowed spillover threshold in kbps. The default is 2000000 Kbps. The valid range is 1 to 2147483647 Kbps.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

DNS - Networking

Set the primary and secondary DNS from this page.

Prerequisites

- Read-Write permission for Networking settings.

Access

- From the Dashboard, go to **Networking > DNS**.
- To override these settings, go to **Networking > Interface**.

Settings

- Enter the primary and secondary DNS addresses, then **Save**.
- Click **Refresh** to view your changes.

Static route settings

Network systems maintain route tables to determine where to forward TCP/IP packets. Routes for outbound traffic are chosen according to the following priorities:

- Direct route
- LLB Link Policy route
- Policy based route
- Static /Dynamic route
- LLB default route

The system evaluates content route rules first, then policy routes, then static routes. The packets are routed to the first route that matches. The LLB route table, therefore, is the one that must include a “default route” to be used when no more specific route has been determined.

Static routes specify the IP address of a next-hop router that's reachable from that network interface. Routers are aware of which IP addresses are reachable through various network pathways, and can forward those packets along pathways capable of reaching the packets’ ultimate destinations. The FortiWAN system itself does not need to know the full route, as long as the routers can pass along the packet.

Note: in FortiWAN implementation, we don't use the standard method (adding a static route with destination 0.0.0.0/0) to add a default route. Instead, the system will generate a default link-load-balance flow-policy to do default routing.

Prerequisites

- Read-Write permission for Networking settings.

Access

- From the Dashboard, go to **Networking > Static Routing**, then select the **Static** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Settings	Guidelines
Destination	Address/mask notation to match the destination IP in the packet header. If there is not a more specific static route defined for a packet's destination IP address, a default route will match the packet, and pass it to a gateway router so that any packet can reach its destination. If you don't define a default route, and if there is a gap in your routes where no route matches a packet's destination IP address, packets passing through the FortiWAN towards those IP addresses will, in effect, be null routed. While this can help to ensure that unintentional traffic can't leave your FortiWAN and therefore can be a type of security measure, the result is that you must modify your routes every time that a new valid destination is added to your network. Otherwise, it will be unreachable. A default route ensures that this kind of locally-caused "destination unreachable" problem does not occur. Enter the address/mask notation to match the destination IP in the packet header. 0.0.0.0/0 or ::/0 is not supported, as FortiWAN uses the LLB default link group as the default route. See Flow Policy on page 160.
Gateway	Enter the IP address of the next-hop router where the FortiWAN system will forward packets for this static route. This router must know how to route packets to the destination IP addresses that you have specified, or forward packets to another router with this information. For a direct Internet connection, this will be the router that forwards traffic towards the Internet, and could belong to your ISP. The gateway must be in the same subnet as the interface used to reach it.
Distance	The default administrative distance is 10, which makes it preferred to OSPF routes that have a default of 110. We recommend you don't change these settings unless your deployment has exceptional requirements.

You can also perform this action using CLI commands. See CLI: Static route settings on page 237.

Configuring policy routes

Network systems maintain route tables to determine where to forward TCP/IP packets. Routes for outbound traffic are chosen according to the following priorities:

- **Link local routes** – Self-traffic uses link local routes.
- **Policy based route** – Configured policy routes have priority over default routes.
- **Static route** – Default routes.
- **LLB Link Policy route** – Configured policy routes have priority over LLB default routes.
- **Dynamic route** – OSPF route have priority over LLB default routes.
- **LLB default route** – lower priority than configured routes.

The system evaluates content route rules first, then policy routes, then static routes. The packets are routed to the first route that matches. The LLB route table, therefore, is the one that must include a “default route” to be used when no more specific route has been determined.

Static routes specify the IP address of a next-hop router that's reachable from that network interface. Routers are aware of which IP addresses are reachable through various network pathways, and can forward those packets along pathways capable of reaching the packets’ ultimate destinations. The FortiWAN system itself does not need to know the full route, as long as the routers can pass along the packet.

Configure at least one static route that points to a router, often a router that’s the gateway to the Internet. You might need to configure multiple static routes if you have multiple gateway routers, redundant ISP links, or other special routing cases.

Prerequisites

- Read-Write permissions for System settings.

Access

- From the Dashboard, go to **Networking > Static Routing**, then select the **Static Routing or Policy Routing** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Settings	Guidelines
Source	Address/mask notation to match the source IP in the packet header. To match any value, either leave it blank or enter 0.0.0.0/32.
Destination	Address/mask notation to match the destination IP in the packet header. To match any value, leave it blank or enter 0.0.0.0/32.
Gateway	IP address of the next-hop router where the FortiWAN system will forward packets for this policy route. This router must know how to route packets to the destination subnet, or forward packets to another router with this information.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

OSPF settings

OSPF (Open Shortest Path First) is described in RFC2328, OSPF Version 2. It is a link-state interior routing protocol. Compared with RIP, OSPF can provide scalable network support and faster convergence times. OSPF is widely used in large networks such as ISP backbone and enterprise networks. FortiWAN supports OSPF version 2.

Prerequisites

- Know how OSPF has been implemented in your network, including the configuration details of the implementation.

Access

- From the Dashboard, go to **Networking > Dynamic Routing**, then select the **OSPF** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Router ID	32-bit number that sets the router-ID of the OSPF process. The router ID uses dotted decimal notation. The router-ID must be an IP address of the router, and it must be unique within the entire OSPF domain to the OSPF speaker.
Default Metric	Enter a number. The default is 10. The valid range is 0 to 16777214.
Distance	Enter a number for the administrative distance. Administrative distance is the measure that routers use to select the best path in a network with multiple routing protocols. The smaller the administrative distance value, the more reliable the routing protocol. The default is 110. The range is 1 to 255.
Redistribute Connected	
Redistribute Connected	Toggle **ON** to enable to redistribute connected routes into OSPF, with the metric type and metric set if specified. Redistributed routes are distributed into OSPF as Type-5 External LSAs into links to areas.
Redistribute Connected Metric	Enter a number. When the OSPF routers redistribute connected routes learned from other protocols, these routes could be distributed with different metric type. OSPF supports two types of external metrics: Type 1 and Type 2. The difference between the two metrics is how OSPF calculates the cost of the route. The default is 10. The valid range is 0 to 167777214.
Redistribute Connected Metric Type	• **1** - Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router.

Setting	Guidelines
	• **2** - Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router.
Redistribute Connected Interface List	Add the items you want to the left side.
Default Information Originate	
Default Information Originate	• **Enable** - Originate an AS-External (type-5) LSA describing a default route into all external routing capable areas of the specified metric and metric type. • **Always** - The default is always advertised, even when there is no default present in the routing table. • **Disable**
Default Information Metric	Enter a number. The default is 10. The valid range is 0 to 167777214.
Default Information Metric Type	• **1** - Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router. • **2** - Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router.
Redistribute Static	
Redistribute Static	Toggle **ON** to redistribute static routes into OSPF, with the metric type and metric set if specified. Redistributed routes are distributed into OSPF as Type-5 External LSAs into links to areas.
Redistribute Static Metric	Enter a number. The default is 10. The valid range is 0 to 167777214.

Setting	Guidelines
Redistribute Static Metric Type	• 1 • 2
Area	
Area	32-bit number that identifies the OSPF area. An OSPF area is a smaller part of the larger OSPF AS. Areas are used to limit the link-state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.
Setting	Guidelines
Authentication	Choose an authentication type: • **None** - Also called null authentication. No authentication is used. In this case the 16-byte Authentication field is not checked, and can be any value. However check summing is still used to locate errors. • **Text** - A simple password is used. The password is a plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication. • **MD5** - Use OSPF cryptographic authentication. A shared secret key is used to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

Setting	Guidelines
Network	
Prefix	Address/mask notation to specify the subnet.
Area	Select an area configuration.
Interface	
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Interace	Select the interface to enable OSPF for it.
Ignore MTU	Enable/disable to ignore the interface MTU. Disabled by default.
Network Type	• Broadcast • Point to Point • Point to Multipoint
Restransmit Interval	Interval for retransmitting Database Description and Link State Request packets. The default is 5 seconds.
Transmit Delay	Increment LSA age by this value when transmitting. The default is 1 second.
Cost	Set link cost for the specified interface. The cost value is set to router-LSA's metric field and used for SPF calculation. The default is 0.
Priority	The router with the highest priority will be more eligible to become Designated Router. Set- ting the value to 0 makes the router ineligible to become Designated Router. The default is 1.
Dead Interval	Number of seconds for RouterDeadInterval timer value used for Wait Timer and Inactivity Timer. This value must be the same for all routers attached to a common network. The default is 40 seconds.
Hello Interval	Number of seconds between hello packets sent on the configured interface. This value must be the same for all routers attached to a common network. The default is 10 seconds.

Setting	Guidelines
Authentication	Choose an authentication type. All OSPF interfaces that want to learn routes from each other must be configured with the same authentication type and password or MD5 key (one match is enough). Options are: • **None** - Also called null authentication. No authentication is used. In this case the 16-byte Authentication field is not checked, and can be any value. However check-summing is still used to locate errors. • **Text** - A simple password is used. The password is a plain text string of characters. The same password is used for all transactions on a network. The main use of this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication. • **MD5** - Use OSPF cryptographic authentication. A shared secret key is used to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.
Text	If using text authentication, enter a password string. Passwords are limited to 8 characters.
MD5	If using MD5 authentication, select an MD5 key object. See Routing – OSPF MD5 key settings on page 91.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Resources

Routing – OSPF MD5 key settings

All OSPF interfaces that want to learn routes from each other must be configured with the same authentication type and password or MD5 key(one match is enough).

Use OSPF cryptographic authentication. A shared secret key is used to authenticate all router traffic on a network. The key is never sent over the network in the clear—a packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Resources > Routing**, then select the **OSPF MD5 Key List** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.
- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Key ID	A number 1-255. Each member key ID must be unique to its member list.
Key	A string of up to 16 characters to be hashed with the cryptographic MD5 hash function.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.

- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** ✖ icon at the end of the row.

Health checks

In server load balancing deployments, the system uses health checks to poll the members of the real server pool to test whether an application is available. You can also configure additional health checks to poll related servers, and you can include results for both in the health check rule. For example, you can configure an HTTP health check test and a RADIUS health check test. In a web application that requires user authentication, the web server is deemed available only if the web server and the related RADIUS server pass the health check.

In link load balancing deployments, the health check can poll either the ISP link group member itself or a “beacon” server that's deployed on the other side of the ISP link. A beacon is an IP address that must be reachable in order for the link to be deemed available. A beacon can be any IP address, such as a main office, core router, or virtual server at another data center.

If you expect a backend server to be unavailable for a long period, such as when it is undergoing hardware repair, it is experiencing extended down time, or when you have removed it from the server farm, you can improve the performance of the FortiWAN system by setting the status of the pool member to Disabled, rather than allowing the system to continue to attempt health checks.

Listed are the predefined health checks. You can also create custom health check objects.

Predefined	Description
LB_HLTHCK_HTTP	Sends a HEAD request to the server port 80. Expects the server to return an HTTP 200.
LB_HLTHCK_ICMP	Pings the server.
LB_HLTHCK_TCP_ECHO	Sends a TCP echo to server port 7. Expects the server to respond with the corresponding TCP echo.
LLB_HLTHCK_ICMP	Pings the destination IP. If the Destination IP is set as default 0.0.0.0, system will ping the default gateway of the link.
LLB_HLTHCK_HOPS	Default hops number is 3. System will ping 8.8.8 by default, with ttl 3 (the default hop number). If the intermediate router responds "ICMP time exceeded in-transit", health check will be successful.

Prerequisites

- Have a good understanding of TCP/IP and knowledge of the services running on your backend servers.
- Know the IP address, port, and configuration details for the applications running on backend servers. For some application protocol checks, you must enter user credentials.
- Read-Write permission for Shared Resources settings.
- After you have configured a health check, you can select it in the SLB server pool, LLB link group, or GLB server configuration.

Access

- From the Dashboard, go to **Resources > Health Check**, then select the **Health Check (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
General	
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Type	Select a health check type: •ICMP •TCP Echo •TCP •HTTP •DNS •RADIUS •SMTP •POP3 •IMAP4 •RADIUS Accounting •FTP •TCP Half Open Connection •TCP SSL •SNMP •SSH •L2 Detection •UDP •SIP •SIP-TCP •SNMP-Custom •RSTP •Hops
Destination Address Type	•IPv4 •IPv6

Setting	Guidelines
Destination IP	IP address to send health check traffic. In server load balancing deployments, if you don't specify an IP address, the real server IP address is used. You might configure IP address for a health check if you are configuring a combination of health checks to poll related servers. In link load balancing deployments, if you don't specify an IP address, the destination IP address is the address of the gateway. You can configure IP address if you want to test connectivity to a beacon on the other side of the gateway, or if you want to test whether service traffic is allowed to pass through the link.
Up Retry	Attempts to retry the health check to confirm server availability. The default is 1.
Down Retry	Attempts to retry the health check to see if a down server has become available. The default is 1. The valid range is 1-10.
Interval	Seconds between each health check. Should be more than the timeout to prevent overlapping health checks. The default is 10.
Timeout	Seconds to wait for a reply before assuming that the health check has failed. The default is 5.
ICMP	
No specific options	Simple ping to test connectivity.
TCP Echo	
No specific options	Simple ping to test connectivity.
TCP / TCP Half Open Connection / UDP	
Port	Listening port number of the backend server. Usually HTTP is 80, FTP is 21, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161.
TCP SSL	

Setting	Guidelines
Port	Listening port number of the backend server. Usually HTTP is 80, FTP is 21, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161.
SSL Ciphers	Default selections are recommended.
Allow SSL Version	Default selections are recommended.
Local Cert	For TCP SSL only. Click the down arrow and select a local SSL Health Check Client certificate from the list menu. The certificate titled "Factory" is the default certificate shipped with your FortiWAN. The rest, if any, are the custom certificates that you have created.
HTTP	
Port	Listening port number of the backend server. Usually HTTP is 80. If testing an HTTP proxy server, Enter the proxy port.
HTTP CONNECT	If the real server pool members are HTTP proxy servers, Enter an HTTP CONNECT option: • **Local CONNECT** - Use HTTP CONNECT to test the tunnel connection through the proxy to the remote server. The member is deemed available if the request returns status code 200 (OK). • **Remote CONNECT** - Use HTTP CONNECT to test both the proxy server response and remote server application availability. If you select this option, you can configure an HTTP request within the tunnel. For example, you can configure an HTTP GET/HEAD request to the specified URL and the expected response. • **No CONNECT** - don't use the HTTP CONNECT method. This option is the default. The HTTP CONNECT option is useful to test the availability of proxy servers only.
Remote Host	If you use HTTP CONNECT to test proxy servers, Enter the remote server IP address.
Remote Port	If you use HTTP CONNECT to test proxy servers, Enter the

Setting	Guidelines
	remote server port.
Method Type	HTTP method for the test traffic: • **HTTP GET** - Send an HTTP GET request to the server. A response to an HTTP GET request includes HTTP headers and HTTP body. • **HTTP HEAD** - Send an HTTP HEAD request. A response to an HTTP HEAD request includes HTTP headers only.
Send String	The request URL, such as /contact.php.
Receive String	A string expected in return when the HTTP GET request is successful.
Status Code	The health check sends an HTTP request to the server. Enter the HTTP status code in the server reply that indicates a successful test. Typically, you use status code 200 (OK). Other status codes indicate errors.
Match Type	What determines a failed health check? • Match String • Match Status • Match All (match both string and status) Not applicable when using HTTP HEAD. HTTP HEAD requests test status code only.
DNS	
Domain Name	The FQDN, such as www.example.com, to use in the DNS A/AAAA record health check.
Address Type	• IPv4 • IPv6
Host IP	IP address that matches the FQDN, indicating a successful health check.
RADIUS / RADIUS Accounting	

Setting	Guidelines
Port	Listening port number of the backend server. Usually RADIUS is 1812 and RADIUS accounting is 1813.
Username	User name of an account on the backend server.
Password	The corresponding password.
Password Type	• **User** - If the backend server does not use CHAP, select this option. • **CHAP** - If the backend server uses CHAP and does not require a secret key, select this option.
SIP / SIP-TCP	
Port	Enter the port number. Valid values range from 0 to 65535.
SIP Request Type	Enter the SIP request type to be used for health checks: • SIP Options • SIP Register
Status Code	The expected response code. If not set, response code 200 is expected. Enter 0 if any reply should indicate the server is available.
SMTP	
Port	Listening port number of the backend server. Usually SMTP is 25.
Domain Name	The FQDN, such as www.example.com, to use in the SMTP HELO request used for health checks. If the response is OK (250), the server is considered as up. If there is error response (501) or no response at all, the server is considered down.
POP3	
Port	Listening port number of the backend server. Usually POP3 is 110.

Setting	Guidelines
Username	User name of an account on the backend server.
Password	The corresponding password.
IMAP4	
Port	Listening port number of the backend server. Usually IMAP4 is 143.
Username	User name of an account on the backend server.
Password	The corresponding password.
Folder	Select an email mailbox to use in the health check. If the mailbox does not exist or isn't accessible, the health check fails. The default is INBOX.
FTP	
Port	Listening port number of the backend server. Usually FTP is 21.
Username	User name of an account on the backend server.
Password	The corresponding password.
File	Enter a file that exists on the backend server. Path is relative to the initial login path. If the file does not exist or isn't accessible, the health check fails.
SNMP	
Port	Listening port number of the backend server. Usually SNMP is 161 or 162.
CPU	Maximum normal CPU usage. If overburdened, the health check fails.
Memory	Maximum normal RAM usage. If overburdened, the health check fails.
Disk	Maximum normal disk usage. If the disk is too full, the health check fails.

Setting	Guidelines
Agent Type	• UCD • Win2000
Community	Must match the SNMP community string set on the backend server. If this does not match, all SNMP health checks fail.
Version	v1 or v2c
CPU Weight	100
Memory Weight	100
Disk Weight	100
SNMP-Custom	
Port	Listening port number of the backend server. Usually SNMP is 161 or 162.
Community	Must match the SNMP community string set on the backend server. If this does not match, all SNMP health checks fail.
Version	V1 or V2C
OID Name	String specifying the OID to query
	Abstract syntax notation (ASN) value type:
SNMP Type	• ASN_INTEGER • ASN_OCTET_STR • ASN_COUNTER • ASN_UINTEGER
SNMP Counter	Enter the value for the evaluation.
SNMP Compare	• Equal • Less • Greater
Name	Enter the name.

Setting	Guidelines
Weight	Enter the weight.
SSH	
Port	Listening port number of the backend server. Usually SSH is 22.
Username	User name for test login.
Password	Corresponding password.
L2 Detection	
No specific options	Link Layer health checker. Sends ARP (IPv4) or NDP (IPv6) packets to test whether a physically connected system is available.
RTSP	
Port	Enter the listening port number. Valid values range from 0 to 65535.
RTSP Method Type	• **RTSP OPTIONS** - returns the request types that the server accepts. • **RTSP DESCRIBE** - includes an RTSP URL (rtsp://…), and the type of reply data that can be handled.
Status Code	200
Hops	
Hops	Default hops number is 3

- Click **Clone** at the right side of column to clone the configuration with a new name.

In SLB deployments, a health check port configuration specifying port 0 acts as a wildcard. The port for health check traffic is imputed from the real server pool member.

In LLB and GLB deployments, specifying port 0 is invalid because there is no associated configuration to impute a proper port. If your health check port configuration specifies port 0, you will not be able to use it in an LLB or GLB configuration.

To disable Health checks in CLI, see CLI Disable Health Check on page 238.

Immediate Health Check

FortiWAN enables you to monitor the health of server in real time directly from your desktop.

Prerequisites

- Have a good understanding of TCP/IP and knowledge of the services running on your backend servers.
- Know the IP address, port, and configuration details for the applications running on backend servers. For some application protocol checks, you must enter user credentials.
- Read-Write permission for Shared Resources settings.
- After you have configured a health check, you can select it in the SLB server pool, LLB link group, or GLB server configuration.

Access

- From the Dashboard, go to **Resources > Health Check**, then select the **Immediate Health Check** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
IP Address	Enter the IP address of the remote server.
Health Check	Select the health check configuration.
Port	Enter the port number, if applicable. This field is only available for health check configurations that require port numbers.

- Click **Start** to perform the health check. The result appears in the **Monitor Information** section.

Address

Create address objects to specify matching source and destination addresses in policies. The following policies use address objects:

- Firewall - see Firewall on page 176
- QoS policies - see Bandwidth on page 169
- Connection limit on page 172
- Link load balance (LLB) - Services on page 160

Workflow

1. Create address objects.
2. Select them to configure address groups or policies. See Address group on page 104

Prerequisites

- Read-Write permission for Shared Resources settings.

For link load balancing, you can also add address objects to address groups, which can then be used in link load balance policies.

Access

- From the Dashboard, go to **Resources > Address**, then select the **Address (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Type	Choose a type from the drop down list.
IPv4/Netmask	
IPv4/Netmask	Enter the IP address.
IP Range	
Address Range	Starting IP address for the range.
To	Ending IP address for the range.
FQDN	
FQDN	Enter the FQDN (fully qualified domain name).
Geography	
Country	Choose an option from the drop down list.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Address group

Configure address groups when you have more than one address object you want to specify in rules that match source or destination addresses. For example, if you subscribe customer 1 and customer 2 to a group of links, then you can create rules that match the customer 1 OR customer 2 address space and load balance the set of gateways assigned to them.

The following policies use address groups:

- Link load balancing policies

Workflow

1. Create address objects. See Address on page 103.
2. Select them to configure address groups or policies.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Resources > Address**, then select the **Address Group** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.
- Click **Add** to open the configuration editor.

Settings	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Member List	
Selected Items	Double-click items to remove members from the group.
Available Items	Double-click items to add members to the group.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

ISP address

Configure ISP addresses to use in policies from this page.

ISP address books contain IP subnet addresses and associated province location settings for ISP links. The following policies use the ISP address book objects:

- GLB data center configuration

The province setting is used in GLB deployments in China to enable location awareness that's province- specific. For example, a user can be directed to a data center in specific location inside the country, such as Beijing or Guangdong, rather than simply China.

Three types of address book entries:

- **Predefined** - Addresses and associated province location settings for China Mobile, China Telecom, and China Unicom. The IP subnet addresses in the predefined address books are not exposed in the user interface. The predefined package is provided to make it easier for you to configure a route when all you know and all you need to know is the name of the ISP that hosts the link.
- **Restored** - Addresses imported from a text file. The IP subnet addresses in the restored address books are not exposed in the user interface. “Restored” addresses can help you rapidly build an ISP address book configuration. “Restored” addresses can help you rapidly build an ISP address book configuration.
- **User-defined** - In the ISP address configuration, you can modify the predefined and restored address books by specifying subnets to add or exclude from them. This gives you flexibility in case you encounter address conflicts or the ISP instructs you to add a subnet address manually.

You can also create new user-defined entries for other ISPs.

In systems with multiple VDOMs, these commands apply to the current VDOM only. In other words, if you configure an exclusion, it is applicable to the current VDOM only; it does not change the predefined address book.

Prerequisites

- Read-Write permission for Shared Resources settings.

Access

- From the Dashboard, go to **Resources > Address**, then select the **ISP Address** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.
  - **Name** - Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
- Click **Restore** to upload a back up file.

  The text file for Restored entries has the following format:

  ```
  #this is a comment line
  ISP name:ABC
  Province:Beijing
  ```

  Managing ISP address books

  ```
  1.1.1.0/24
  Province:Unknown 2.2.0.0 255.255.0.0
  #this is a comment line too 3.3.3.3/32
  ISP name:DEF Province:Shanghai 4.4.4.0 255.255.255.0
  5.5.0.0/16
  ```

  - Click **Back Up** to create a back up file.
  - Click **Clean** to erase entries that were imported from the text file. The clean operation does not affect the predefined addresses or user-configured entries. If a restored entry has user-configured elements (for example, an exclude list), the clean operation clears the addresses but preserves the configuration and converts it to a user- defined type.
  - Click **Inquire**, then enter a specific IP address to see whether an IP address belongs to any of the address books. If an address can be found in more than one address book, the results are returned in the following priority:
    1. User-defined
    2. Restored
    3. Predefined
  - To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
  - To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
  - Click **Clone** at the right side of column to clone the configuration with a new name.

Applications

FortiWAN provides more than two dozen predefined applications, and allows you to create your application objects. Application objects are an important part of the following policy configurations:

- Firewall policies
- QoS policies

- Connection limit policies
- Link load balancing policies

Workflow

1. Create application objects.
2. Select them when you configure application groups or policies.

Prerequisites

- Read-Write permission for Shared Resources settings.

For link load-balancing, you can add application objects to application groups; then use the application groups in LLB policies.

Access

- From the Dashboard, go to **Resources > Application**, then select the **Application (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
APP_ID	There are two types of APP-ID. 1) predefined, and 2) user-defined. If you are setting a user-defined APP-ID, valid characters are 0-9. No spaces or letters. The custom application range is 1000001 - 2000000.
Application Type	• **Port Based** -define a custom application by protocol type and port number • **Internet Service** - define a custom application by protocol number plus "IP Range", "IP netmask" or "FQDN"
Protocol Type	Select one of the following: • IP (default) • ICMP • TCP • UDP • TCP-and-UDP • SCTP
Protocol	This applies when Protocol Type is to set to **IP**. In that case, it displays the protocol number without port. The default value is 1
These options become available when TCP, UDP, SCTP, or TCP-AND-UDP is selected.	
Minimum Destination Port	Enter a port number. The default value is 1. The valid range is 1 to 65535.
Maximum Destination Port	Enter a port number. The default value is 65535. The valid range is 1 to 65535.
Specify Source Port	Toggle **ON** to specify the source port.
Source Port Min	This applies only when the specify source port is enabled.

Setting	Guidelines
	The default value is 1.
Source Port Max	This applies only when the specify source port is enabled. The default value is 65535.

- Click **Clone** at the right side of column to clone the configuration with a new name.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page.

Application groups

Configure application groups when you have more than one application and you want to specify in a rule a match to the application. You can group all Web applications and group all mail applications, for example, if you want to have rules that treat those as groups.

The following policies use application groups:

- Link load balancing policies

Workflow

1. Create application objects. See Applications on page 107.
2. Configure application group objects.
3. Select the application groups when you configure your policies.

Prerequisites

- Read-Write permission for Shared Resources settings.

Access

- From the Dashboard, go to **Resources > Application**, then select the **Application Group** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Member List	
Selected Items	Double-click items to remove members from the group.
Available Items	Double-click items to add members to the group.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Schedules and schedule groups

Create schedule objects to use in link load balancing policies. A policy rule can be time-bound: one time, daily, weekly, or monthly.

Workflow

1. Create a schedule object.
2. Select the schedule when you configure the link policy.

Prerequisites

- Read-Write permission for Shared Resources settings.

Access

- From the Dashboard, go to **Resources > Schedule**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

1. Click **Add** to open the configuration editor.
2. Enter a name, then click **Save**.

   Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.

   The schedule appears in the list.
3. Double-click the item from the list, or click the **Pencil** icon at the end of the row.

   The Member pane appears.
4. Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Type	• One Time • Daily • Weekly
Start Date	YYYY/MM/DD
End Date	YYYY/MM/DD
Start Time	HH:MM
End Time	HH:MM

5. Continue adding members as needed.
   - To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
   - To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
   - Click **Clone** at the right side of column to clone the configuration with a new name.

Link load balance (LLB) - Resources

The system has a configuration framework that enables granular link load balancing rules.

The following shows the configuration objects used in the LLB configuration and the order in which you create them. A LLB flow policy specifies the source-destination-application matches to which the policy applies. Apply a link policy to a underlay link or an overlay link.

Figure 9 - Link Load Balancing configuration

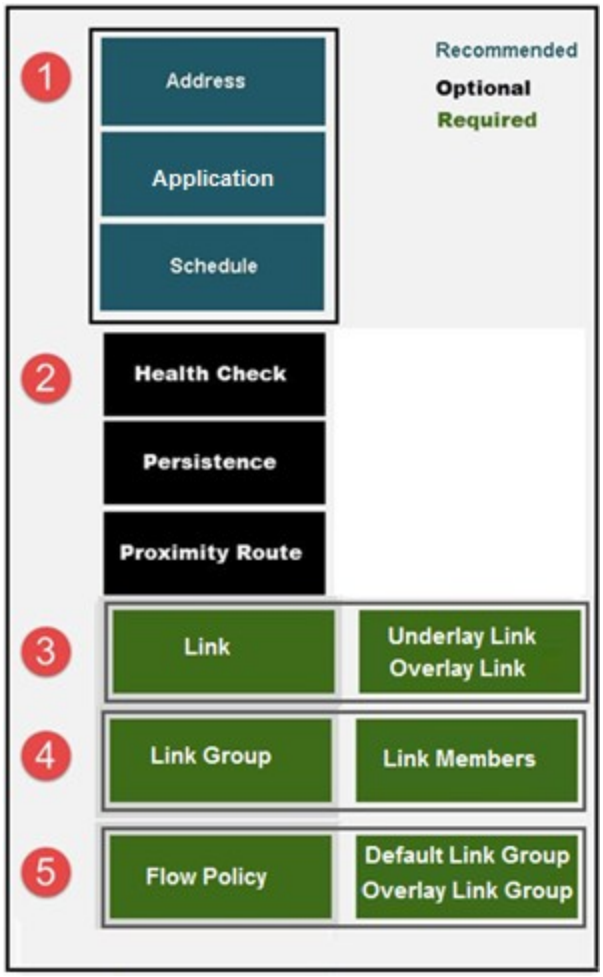

- The granular configuration of the underlay links configuration includes health checks and bandwidth thresholds.
- The granular configuration of Overlay links includes spillover In/ Out/ Total thresholds, but don't use health check configuration objects for inbound or outbound bandwidth.
- The granular configuration of link groups includes load balancing methods, persistence rules, and proximity routes.

Workflow

1. Add address, address group, application, app group, and schedule group configuration objects that can be used to match traffic to link policy rules. This step is recommended. If your policy does not use match criteria, it will not have granularity. See Link load balance (LLB) - Services on page 160.
2. Configure optional features. If you want to use health check rules, configure them before you configure the gateway links. If you want to use persistence rules or proximity routes, configure them before you configure a link group. See Links - underlay or overlay on page 76 and Link group on page 116.
3. Configure IPsec Tunnels. See IPsec tunnels on page 71.
4. Configure LLB link for Underlay and Overlay. See below.
5. Configure link groups. See Link group on page 116.
6. Configure the flow policy. When you configure a link policy, you set the source-destination-application matching tuple for your link groups or virtual tunnels. See Flow Policy on page 160.

LLB overlay links

The data path between a user's computer from Hub device and a private network from Branch device through a VPN is referred to as a tunnel. This IPSec tunnel is also called Overlay Links. Like a physical tunnel, the data path is accessible only at both ends.

FortiWAN restricts that when all devices belong to the same group, only one device is allowed as the HUB node role, and at the same time, multiple branch nodes are allowed. All the branch node devices access the HUB node device to connect and communicate through the VPN tunnel.

When there are multiple Group communication, you need to establish VPNs between Hub nodes of different Groups. The client packet is sent by the Branch node of the client to the Hub node of the Local Group. The packet passes through the VPN tunnel between the HUB and the HUB to reach the peer branch node, and finally descrypts and reaches the destination address.

Encapsulation makes this possible. IPsec packets pass from one end of the tunnel to the other and contain data packets that are exchanged between the local user and the remote private network. Encryption of the data packets ensures that any third-party who intercepts the IPsec packets can not access the data.

An IPsec tunnel is a good choice when you want to load balance traffic from applications that embed the source address in the packet payload, like VPN and VoIP traffic. Such traffic can be difficult to load balance using traditional LLB methods. IPsec tunnels enable reliable, site-to-site connectivity using IPSec tunnel. The local FortiWAN appliance encapsulates traffic so that it can be routed according to your link flow policy rules. The link flow policy rules use LLB techniques to identify the best available route among a group of links. If one of the links breaks down, the traffic

can be rerouted through another link in the tunnel group. When traffic egresses the remote FortiWAN appliance, it is decapsulated and the original source and destination IP addresses are restored.

The following shows the same network deployed with FortiWAN appliances. The LLB underlay and overlay link policy load balances traffic among more affordable ADSL links.

Figure 10 - LLB overlay link

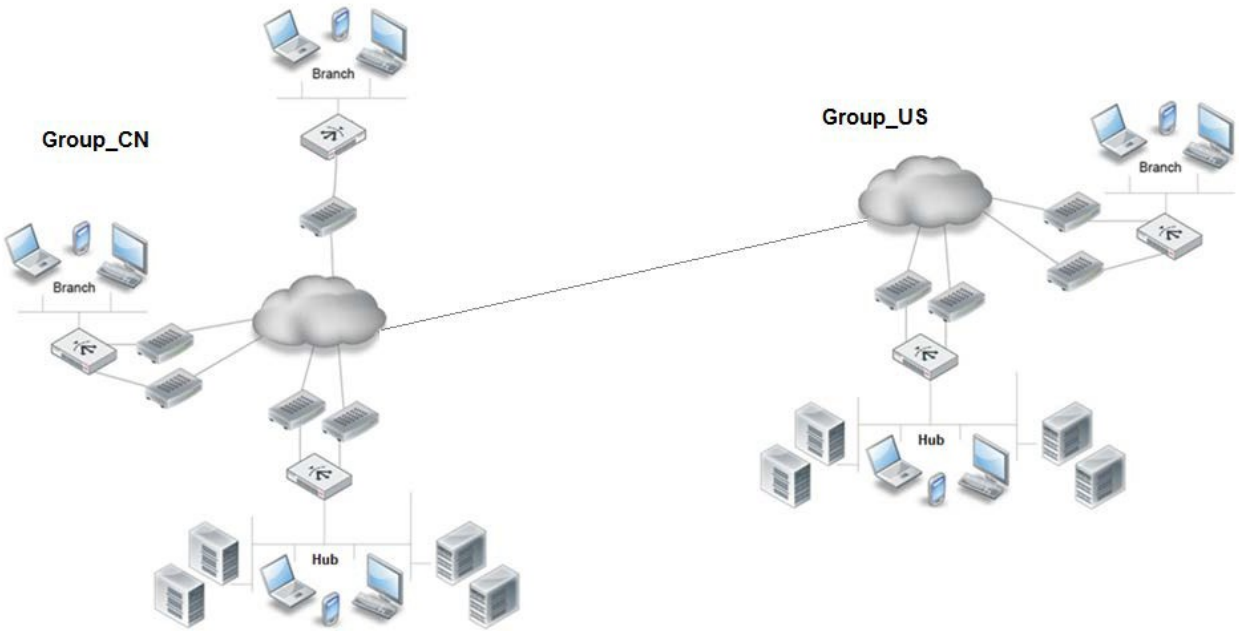

The FortiWAN system evaluates traffic to determine the routing rules to apply. With regard to link load balancing, the system evaluates rules in the following order and applies the first match:

1. Direct route
2. LLB route
3. Policy route
4. Static/Dynamic route
5. LLB default

LLB underlay links

The LLB underlay links option is useful for ISP links. It enables you to configure multiple ISP links that are possible routes for the traffic. The link type FortiWAN supports both static IP and dynamic type. The dynamic link includes DHCP and PPPoE to obtain IP.

The following shows an example topology when FortiWAN is deployed to support LLB link.

Figure 11 - LLB underlay links

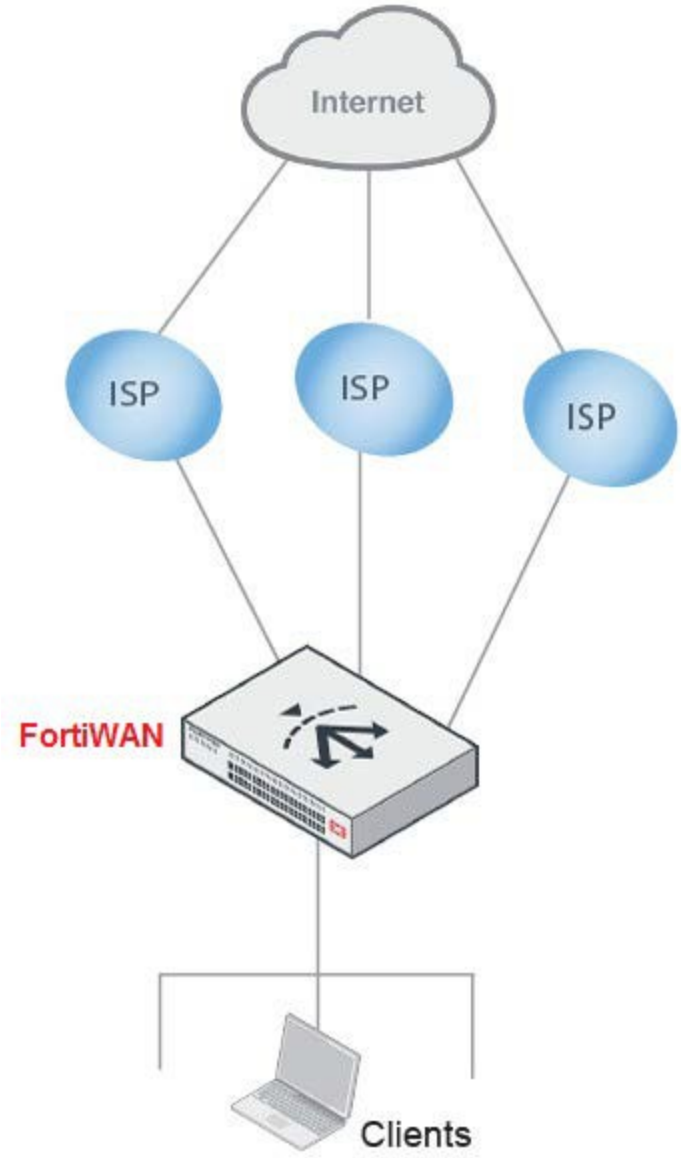

Link group

Link groups include Underlay link group and Overlay link group. Grouping links reduces the risk of outages and provisions additional bandwidth to relieve potential traffic congestion.

The link group configuration specifies the load balancing algorithm. You can enable LLB options, such as persistence rules and proximity routes.

Prerequisites

- Configure gateway links and persistence rules, so you can select them in the link group configuration.
- Read-Write permission for Resources module settings.

After you have configured a link group configuration object, you can select it in the flow policy configuration.

Access

- From the Dashboard, go to **Resources > Link Load Balance**, then select the **Link Group** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Address Type	IPv4 is selected by default, and can't be changed.
Route Method	• **Weighted Round Robin** - Dispatches new connections to link members using a weighted round-robin method. • **Upstream-traffic** - Dispatches new connections to the link member with the least outbound traffic. • **Downstream traffic** - Dispatches new connections to the link member with the least inbound traffic. • **Total traffic** - Dispatches new connections to the link member with the least total traffic (that's, inbound plus outbond). • **Spillover downstream traffic** - Spillover list based on inbound traffic. • **Least RTT** - The fastest round trip time.
Persistence	(optional) Select a persistence configuration from the drop down list.
Proximity Route	Toggle **ON** to use the proximity route logic and configuration when determining routes.
Link Member	
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Link	Select a link configuration object.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1 to 255. All load balancing methods consider weight, except spillover, which uses its own priority configuration. Servers are dispatched requests proportional to their weight, relative to the sum of all weights.

Setting	Guidelines
	The following example shows the effect of weight on WRR: • Sever A, Weight 2; Server B, Weight 1: Requests are sent AABAAB. • Sever A, Weight 3; Server B, Weight 2: Requests are sent AABAB. For other methods, weight functions as a tie-breaker. For example, with the Least Connection algorithm, requests are sent to the server with the least connections. If the number of connections is equal, the request is sent to the server with the greater weight. For example: • Server A, Weight 1, 1 connection • Server B, Weight 2, 1 connection The next request is sent to Server B.
Spillover Priority	Assigns a priority to the link when using a spillover load balancing method. Higher values have greater priority. When a spillover method is enabled, the system dispatches new connections to the link that has the greatest spillover priority until its threshold is exceeded; then it dispatches new connections to the link with the next greatest priority until its threshold is exceeded, and so on. If multiple links in a link group have the same spillover priority, the system dispatches new connections among those links according to round robin. The default is 0. The valid range is 0-9.
Status	Toggle **ON** to consider this group available for new traffic.
Backup	Toggle **ON** to back up this configuration.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Proximity route

The proximity route feature enables you to associate link groups with efficient routes. Proximity routes can improve user experience over the WAN because traffic is routed over fast routes.

You can use either or both of these methods:

- **Static Table** - You specify the gateways to use for traffic on destination networks.
- **Dynamic Detection** - The system polls the network for efficient routes. The algorithm selects a gateway based on latency.

If you configure both, the system checks the static table first for a matching route and, if any, uses it. If there is no matching static route, the system uses dynamic detection.

Prerequisites

- Read-Write permission for Resources module settings.

Access

- From the Dashboard, go to **Resources > Link Load Balance**, then select the **Proximity Route** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Choose the **Mode**, then enter your options.

Setting	Guidelines
Mode	• **Static Table First** - Consult the static table first. If no match, use dynamic detection. • **Static Table Only** - Use the static table; don't use dynamic detection. • **Dynamic Detect Only** - Use dynamic detection; don't use the static table. • **Disable** - don't use the proximity route configuration.
Static Table	
Type	• **ISP** - Use an ISP address object. • **Subnet** - Enter an IP netmask manually. Routes that are specified manually have priority over ISP address object entries.
ISP Name	When using the **ISP** configuration type, select an ISP address book configuration object. If an address exists in multiple ISP address books, the route entries have priority as follows: 1. User-defined entries. 2. Entries from an address book that has been imported. 3. Entries from the predefined address book (default for the firmware image).
IP Subnet	When using the **Subnet** configuration type, specify a destination IP address and netmask.
Gateway	Select a gateway configuration object. The gateway must be able to route packets to the destination IP address that you have specified.
Dynamic Detect	
Protocol	• **ICMP** - Use ICMP to detect routes. Calculate proximity by the smaller RTT. • **ICMP and TCP** - Some hosts don't respond to ICMP

Setting	Guidelines
	requests. Enter this option to use both ICMP and TCP to detect routes and RTT. For TCP detection, port 7 (TCP echo) is used. A connection refused or connection reset by the destination is treated as successful detection.
Aging Period	Enter the number of seconds to detect. When the timer expires, a new detect is triggered and the cache is deleted. The default is 86,400 seconds (24 hours). The range is 60 to 2592000 seconds.
Retry Number	Enter the number of allowable retries. The default is 3. The range is 1 to 10.
Retry Interval	Enter the time between retries, in seconds. The default is 3. The range is 1 to 3600 seconds.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Traffic Shaper

Configure a traffic shaper before you configure a bandwidth policy.

Traffic shaping optimizes performance, improves latency, or increases usable bandwidth for some kinds of packets by delaying other kinds.

Workflow

1. Configure a traffic shaper for bandwidth management.
2. Configure a bandwidth policy. See Bandwidth on page 169.

Prerequisites

- Have Read-Write permission for Link Load Balance settings.

Access

- From the Dashboard, go to **Resources > Link Load Balance**, then select the **Traffic Shaper** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.
- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name..
Maximum Bandwidth	Maximum bandwidth rate. Enter a number and a unit abbreviation. Example: 10G,20M,30K,40 unit: G|M|K|b(default)
Guaranteed Bandwidth	Guaranteed bandwidth rate. Enter a number and a unit abbreviation. Example: 10G,20M,30K,40 unit: G|M|K|b(default)
Priority	Enter the priority for the traffic shaper. • High • Medium • Low

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Persistence - real server

Use this page to configure source address affinity and a timeout for GSLB persistence. You enable persistence per host in the GSLB host configuration.

If the DNS query is for a host that has persistence enabled, the DNS server replies with an answer that has the virtual server IP addresses listed in the order determined by the GSLB proximity algorithms,

and the client source IP address (for example 192.168.1.100) is recorded in the persistence table. If source address affinity is set to 24 bits, subsequent queries for the host from the 192.168.1.0/24 network are sent an answer with the virtual servers listed in the same order (unless a server becomes unavailable and is therefore omitted from the answer).

Persistence is required for applications that include transactions across multiple hosts, so the persistence table is also used for queries for other hosts with the same domain. For example, a transaction on a banking application might include connections to login.bank.com and transfer.bank.com. To support persistence in these cases, the GSLB persistence lookup accounts for domain as well. The first query for login.bank.com creates a mapping for the source address network 192.168.1.0/24 and the domain bank.com. When the DNS server receives subsequent requests, it consults the persistence table for a source network match, then a domain match and a host name match. In this example, as long as you have created host configurations for both login.bank.com and transfer.bank.com, and persistence is enabled for each, the persistence table can be used to ensure the DNS answers to queries from the same network list the resource records in the same order.

Prerequisites

- Read-Write permission for Global Load Balance settings.

Access

- From the Dashboard, go to **Resources > Link Load Balance**, then select the **Persistence** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Type	Choose from the drop-down list.
Source-Destination Pair	
Timeout	Enter the number of seconds to timeout. The default is 300 seconds. The valid range is from 1 to 86400 seconds.
Source-Destination Address	
Timeout	Enter the number of seconds to timeout. The default is 300 seconds. The valid range is from 1 to 86400 seconds.
Source IPv4 Netmask Bits	Enter the number of bits in a subnet mask to specify a network segment that follows the persistence rule. The default is 32. The valid range is from 1 to 32.
Destination IPv4 Netmask Bits	The number of bits in a subnet mask to specify a network segment that follows the persistence rule. For example, if you set this to 24, and the system chooses a particular gateway router for destination IP 192.168.1.100, the system selects that same gateway for traffic to all destination IPs in subnet 192.168.1.0/24. The default is 32. The valid range is from 1 to 32.
Source Address	
Timeout	Enter the number of seconds to timeout. The default is 300 seconds. The valid range is from 1 to 86400 seconds.
Source IPv4 Netmask Bits	Enter the number of bits in a subnet mask to specify a network segment that follows the persistence rule.

Setting	Guidelines
	For example, if you set this to 24, and the system chooses a particular gateway router for client IP 192.168.1.100, the system selects that same gateway for subsequent client requests when the subsequent client belongs to subnet 192.168.1.0/24. The default is 32. The valid range is from 1 to 32.
Destination Address	
Timeout	Enter the number of seconds to timeout. The default is 300 seconds. The valid range is from 1 to 86400 seconds.
Destination IPv4 Netmask Bits	Enter the number of bits in a subnet mask to specify a network segment that follows the persistence rule. The default is 32. The valid range is from 1 to 32.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row. You can't modify a default persistence rule.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row. You can't delete a default persistence rule.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Virtual Servers

These topics describe how to set up resources needed to create virtual servers.

Profiles

An application profile is a configuration object that defines how you want the FortiADC virtual server to handle traffic for specific protocols.

The following lists usage by load balancing profile type, including compatible virtual server types, load balancing methods, persistence methods, and content route types.

Profile	Usage	VS Type	LB Methods	Persistence
FTP	Use with FTP servers.	Layer 4	Round Robin, Least	Source Address, Source Address

Profile	Usage	VS Type	LB Methods	Persistence
			Connections, Fastest Response	Hash
TCP	Use for other TCP protocols.	Layer 4, Layer 2	Round Robin, Least Connections, Fastest Response	Source Address, Source Address Hash
UDP	Use for other UDP protocols.	Layer 4, Layer 2	Round Robin, Least Connections, Fastest Response	Source Address, Source Address Hash

Listed are the default values of the predefined profiles. All values in the predefined profiles are view-only, and cannot be modified. You can select predefined profiles in the virtual server configuration, or you can create user-defined profiles.

Setting	Guidelines
LB_PROF_TCP	Timeout TCP Session - 100 Timeout TCP Session after FIN - 100
LB_PROF_UDP	Timeout UDP Session - 100
LB_PROF_FTP	Timeout TCP Session - 100 Timeout TCP Session after FIN - 100

Prerequisites

- Create configuration objects for certificates, caching, and compression if you want the profile to use them.
- Read-Write permission for Load Balance settings.

Access

- From the Dashboard, go to **Resources > Virtual Server**, then select the **Profile (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
TCP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.
UDP	
Timeout UDP Session	Client-side session timeout. The default is 100 seconds. The valid range is 1 to 86,400.
FTP	
Timeout TCP Session	Client-side timeout for connections where the client has not sent a FIN signal, but the connection has been idle. The default is 100 seconds. The valid range is 1 to 86,400.
Timeout TCP Session after FIN	Client-side connection timeout. The default is 100 seconds. The valid range is 1 to 86,400.

- Click **Clone** at the right side of column to clone the configuration with a new name.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Persistence - virtual server

Persistence rules identify traffic that should not be load balanced, but instead forwarded to the same backend server that has seen requests from that source before. Typically, you configure persistence

rules to support server transactions that depend on an established client-server session, like e-commerce transactions or SIP voice calls.

The system maintains persistence session tables to map client traffic to back end servers based on the session attribute specified by the persistence rule.

The persistence table is evaluated before load balancing rules. If the packets received by the SLBADC match an entry in the persistence session table, the packets are forwarded to the server that established the connection, and load balancing rules are not applicable.

Typical source-address persistence rule have a timeout that you can specified. When the time that has elapsed since the system last received a request from the client IP address is greater than the timeout, the system does not use the mapping table to forward the request. Instead, it again selects the server using the method specified in the virtual server configuration. Source address hash rule has a timeout built into the hash algorithm so that you do not need to specify it.

Following lists the predefined persistence rules. You can get started with these commonly used persistence methods or create custom objects.

Setting	Guidelines
LB_PERSIS_SIP	Persistence based on source IP address or subnet.
LB_PERSIS_ CONSISTENT_SIP	Persistence based on a hash of source IP address.

Prerequisites

- Have a good understanding and knowledge of the applications that require persistent sessions and the methods that can be used to identify application sessions.
- Read-Write permission for Load Balance settings.

After you have configured a persistence rule, you can select it in the virtual server configuration.

Access

- From the Dashboard, go to **Resource > Virtual Server**, then select the **Persistence** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Type	Choose from the drop-down list.
Source Address	
Source Address	Persistence is based on source IP address.
Timeout	Enter the number of seconds to timeout for an inactive persistence session table entry. The default is 300 seconds. The valid range is 1 to 86,400.
Subnet Mask Bits (IPv4)	Number of bits in a subnet mask to specify a network segment that should following the persistence rule. For example, if IPv4 mask bits is set to 24, and the back end server A responds to a client with the source IP 192.168.1.100, server A also responds to all clients from subnet 192.168.1.0/24. The default is 32. The valid range is from 1 to 32.
Subnet Mask Bits (IPv6)	Number of bits in a subnet mask to specify a network segment that should following the persistence rule. The default is 128. The valid range is from 1 to 128.
Match Across Virtual Servers	Enable so clients continue to access the same back end server through different virtual servers for the duration of a session. For example, a client session with a vSphere 6.0 Platform Services Controller (PSC) has connections on the following ports: 443, 389, 636, 2012, 2014, 2020. A FortiWAN deployment to load balance a cluster of vSphere PSCs includes Layer 4 virtual server configurations for each of these ports. To ensure a client's connections for a session go to the same back end real server: • Create a persistence object based on Source Address affinity and select the Match Across Servers option. • Select this persistence object in each of the Layer 4 virtual servers configured to load balance the vSphere PSC pool.

Setting	Guidelines
	• Select the same real server pool object in each of the Layer 4 virtual servers configured to load balance the vSphere PSC pool. When these options are enabled, FortiWAN dispatches the initial connection to a real server destination (for example, RS1) based on the virtual server's load balancing method, and the persistence object is noted in the connection table. Subsequent connection attempts with the same source IP address to any FortiWAN virtual server that has this persistence object and real server pool are dispatched to RS1, as long as the session is active. **Note**: In the Layer 4 virtual server configuration, you specify a packet forwarding method. You can use Source Address persistence with Match Across Servers with any combination of Direct Routing, DNAT, and Full NAT packet forwarding methods. However, with NAT46 and NAT64 packet forwarding methods, the source address type is different from the real server address type. To use Match Across Servers with NAT46 or NAT64, all virtual servers for the application must be configured with the same packet forwarding method: all NAT46 or all NAT64.
Source Address Hash	
Source Address Hash	Persistence is based on a hash of the IP address of the client making an initial request.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row. You can't modify a default persistence rule.
- Click **Clone** at the right side of column to clone the configuration with a new name.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row. You can't delete a default persistence rule.

Real server

Configure real servers on this page.

Real servers are the hosts providing service in the backend. Virtual servers forward service requests to the real servers with specified load balancing method.

Prerequisites

- The IP address of the real servers.
- Read-Write permission for Load Balance settings.

After you have configured a real server, you can select it in the real server pool configuration. See Real server pool on page 133.

Access

- From the Dashboard, go to **Resources > Virtual Server**, then select the **Real Server** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.This is the name of the real server.
Status	Enable/Disable the real server. Virtual server will not forward service requests to disabled real server.
Address	IPv4 address of the real server.
Address6	IPv6 address of the real server.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Real server pool

Configure real server pools from this page.

Server pools are groups of real servers that host the applications that you load balance.

Workflow

- Create a server pool object.
- Add members.

Prerequisites

- Have a good understanding and knowledge of the backend server boot behavior, for example, how many seconds it takes to “warm up” after a restart before it can process traffic.
- The IP address and port of the applications.
- If you want to select user-defined health checks, you must create them before creating the pool configuration.
- Read-Write permission for Load Balance settings.

After you have configured a real server pool, you can select it in the virtual server configuration.

Access

- From the Dashboard, go to **Resources > Virtual Server**, then select the **Real Server Pool** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Real Server Pool	
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. Reference this name in the virtual server configuration.
Address Type	IPv4/IPv6
Health Check	Enable health checking for the pool. You can override this for individual servers in the pool.
Health Check Relationship	• **AND** - All of the selected health checks must pass for the server to the considered available. • **OR** - One of the selected health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects.
Member (Save the above configurations so that you can continue the members)	
Status	• **Enable** - The server can receive new sessions. • **Disable** - The server does not receive new sessions and closes any current sessions as soon as possible. • **Maintain** - The server does not receive new sessions but maintains any current connections.
Real Server	Select the predefined real server.
Port	backend server listening port number. Usually HTTP is 80, HTTPS is 443, FTP is 21, SMTP is 25, DNS is 53, POP3 is 110, IMAP4 is 143, RADIUS is 1812, and SNMP is 161. **Tip**: The system handles port 0 as a “wild card” port. When configured to use port 0, the system uses the destination port from the client request. For example, if you specify 0, and the destination port in the client request is 50000, the traffic is forwarded to port 50000. See “Example: Using port ranges and the port 0 configuration”.

Setting	Guidelines
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1 to 256. All load balancing methods consider weight. Servers are dispatched requests proportional to their weight, relative to the sum of all weights. The following example shows the effect of weight on Round Robin: • Sever A, Weight 2; Server B, Weight 1: Requests are sent AABAAB. • Sever A, Weight 3; Server B, Weight 2: Requests are sent AABAB. For other methods, weight functions as a tie-breaker. For example, with the Least Connection algorithm, requests are sent to the server with the least connections. If the number of connections is equal, the request is sent to the server with the greater weight. For example: • Server A, Weight 1, 1 connection • Server B, Weight 2, 1 connection The next request is sent to Server B.
Health Check Inherit	Use the pool's health check settings. Disable to override those settings by selecting a different health check to use with this individual backend server.
Health Check	Toggle **ON** to enable health checking for this server.
Health Check Relationship	• **AND** - All of the selected health checks must pass for the server to the considered available. • **OR** - One of the selected health checks must pass for the server to be considered available.
Health Check List	Select one or more health check configuration objects. Shift-click to select multiple objects.

To configure a pool

1. Go to **Resource > Virtual Server,** and select the **Real Server Pool** tab.
2. Click **Add** to open the configuration editor.
3. Enter your information into the fields and add real server members.

Example: Using port ranges and the port 0 configuration

In some deployments, it is advantageous to support listening port ranges for client requests. For example, data centers or web hosting companies sometimes use port numbers to identify their customers. Client A sends requests to port 50000, client B to port 50001, client C to port 50002, and so on.

To support this scenario:

- On the real servers, configure the listening ports and port ranges according to your requirements.
- On the FortiWAN, when you configure the real server pool member, specify port 0 for the port. The system handles port 0 as a “wildcard” port. When configured to use port 0, the system uses the destination port from the client request. For example, if you specify 0, and the destination port in the client request is 50000, the traffic is forwarded to port 50000.
- When you configure the virtual server, specify a listening port and port range. The port range is like an offset. If the specified port is 50000 and the port range is 10, the virtual server listens on ports 50000-50009.

Using Source (NAT) pools

Use the **Source Pool** page to create configuration objects for source IP addresses used for NAT in Layer 4 virtual server configurations.

In a Layer 4 virtual server configuration, you select a “packet forwarding method” that includes the following network address translation (NAT) options:

- **Direct Routing** - Does not rewrite source or destination IP addresses.
- **NAT** - Rewrites the destination IP address (DNAT) for packets before it forwards them.
- **Full NAT** - Rewrites both the source and destination IP addresses. Use for standard NAT, when client and server IP addresses are all IPv4 or all IPv6.
- **NAT46** - Rewrites both the source and destination IP addresses. Use for NAT 46, when client IP addresses are IPv4 and server IP addresses are IPv6.

- **NAT64** - Rewrites both the source and destination IP addresses. Use for NAT 64, when client IP addresses are IPv6 and server IP addresses are IPv4.

Prerequisites

- You must have a good understanding of NAT. You must know the address ranges your network has provisioned for NAT.
- Configure the backend servers to use the FortiWAN address as the default gateway so that server responses are also rewritten by the NAT module.
- Read-Write permission for Load Balance settings.

After you have configured a source pool IP address range configuration object, you can select it in the virtual server configuration. You can assign a virtual server multiple source pools (with the same or different associated source pool interfaces).

Access

- From the Dashboard, go to **Monitor > System**, then select the **tab name** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

To configure a source pool

1. Go to **Resource > Virtual Server**, and click the **NAT Pool** tab.
2. Click **Add** to open the configuration editor.
3. Enter your information into the fields.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. Reference this name in the virtual server configuration.
Address Type	IPv4/ IPv6

Setting	Guidelines
Address Range	The first address in the address pool.
To	The last address in the address pool.
Interface	Interface to receive responses from the backend server. The interface used for the initial client traffic is determined by the virtual server configuration.

Example: DNAT

The NAT module rewrites only the destination IP address. Therefore, if you configure destination NAT, you do not need to configure a source pool. In this DNAT example, the destination IP address in the packets it receives from the client request is the IP address of the virtual server—192.168.1.101. The NAT module translates this address to the address of the real server selected by the load balancer—in this example, 192.168.2.1. The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic..

Figure 12 - Destination NAT

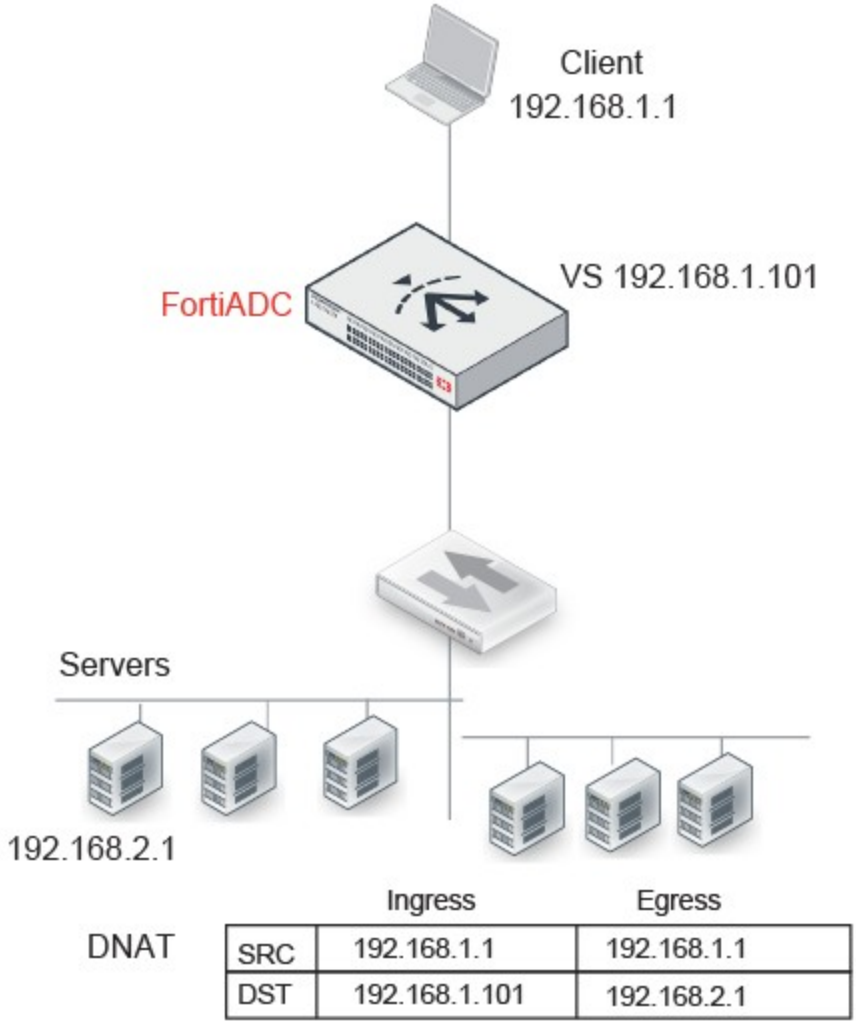

		Ingress	Egress
DNAT	SRC	192.168.1.1	192.168.1.1
	DST	192.168.1.101	192.168.2.1

Example: FullNAT

The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the next available address in the source pool—in this example, 192.168.2.101. It translates the destination IP address to the address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic..

Figure 13 - Full NAT

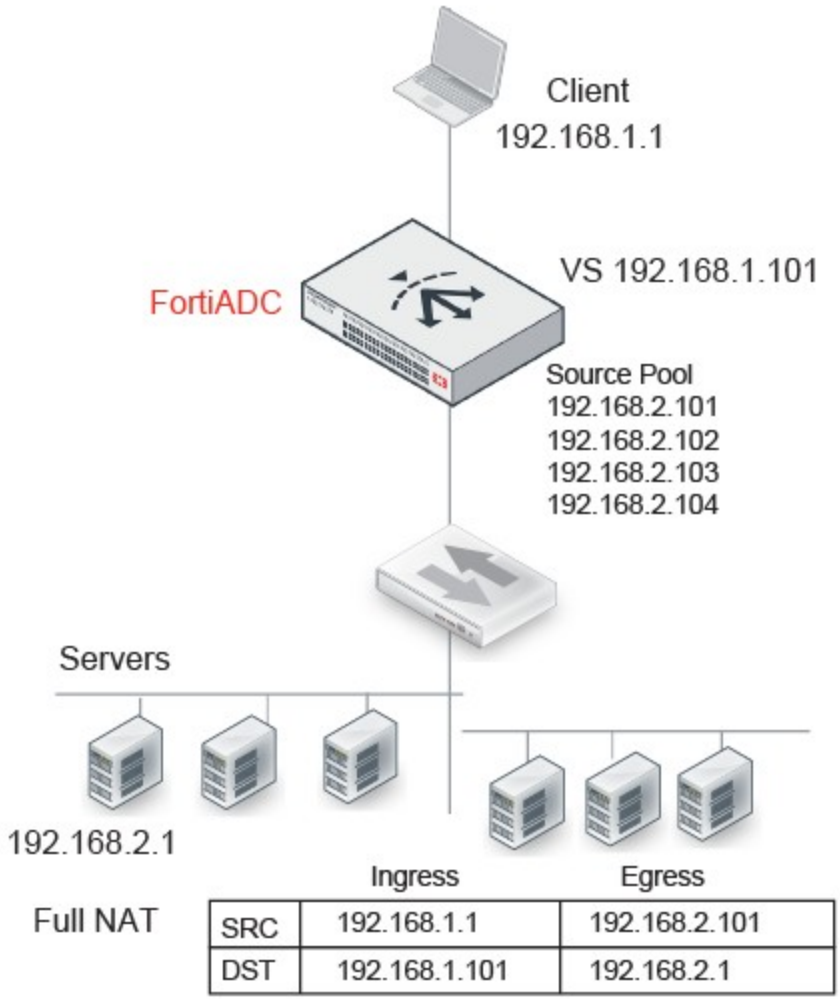

Example: NAT46 (Layer 4 virtual servers)

The IPv6 client connects to the virtual server IPv4 address. The source IP / destination IP pair in the packets received is SRC 192.168.1.1 / DST 192.168.1.101. The NAT module translates the source IP address to the next available IPv6 address in the source pool—in this example, 2002::2:1001. It translates the destination IP address to the IPv6 address of the real server selected by the load balancer—in this example, 2002::2:1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic...

Figure 14 - Full NAT with NAT46

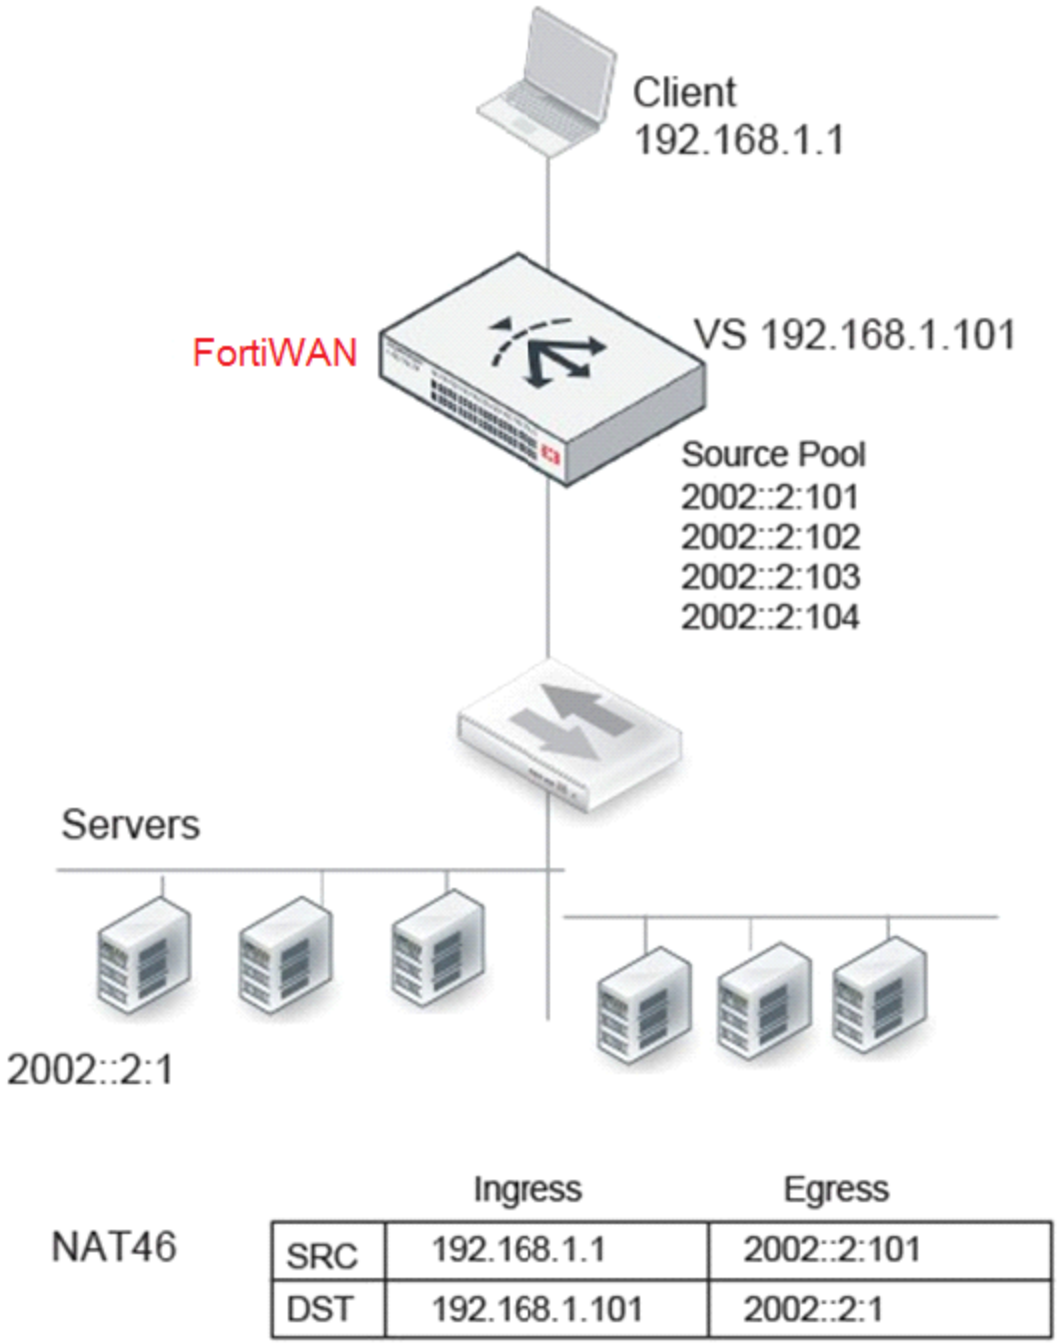

NAT46

	Ingress	Egress
SRC	192.168.1.1	2002::2:101
DST	192.168.1.101	2002::2:1

Features	Notes
Profile	Not Supported: FTP
ICMP	ICMP traffic is dropped.

Example: NAT64 (Layer 4 virtual servers)

The IPv6 client connects to the virtual server IPv6 address. The source IP / destination IP pair in the packets received is SRC 2001::1:1 / DST 2001::1:101. The NAT module translates the source IP address to the next available IPv4 address in the source pool—in this example, 192.168.2.101. It translates the destination IP address to the IPv4 address of the real server selected by the load balancer—in this example, 192.168.2.1.

The system maintains this NAT table and performs the inverse translation when it receives the server-to-client traffic...

Figure 15 - Full NAT with NAT64

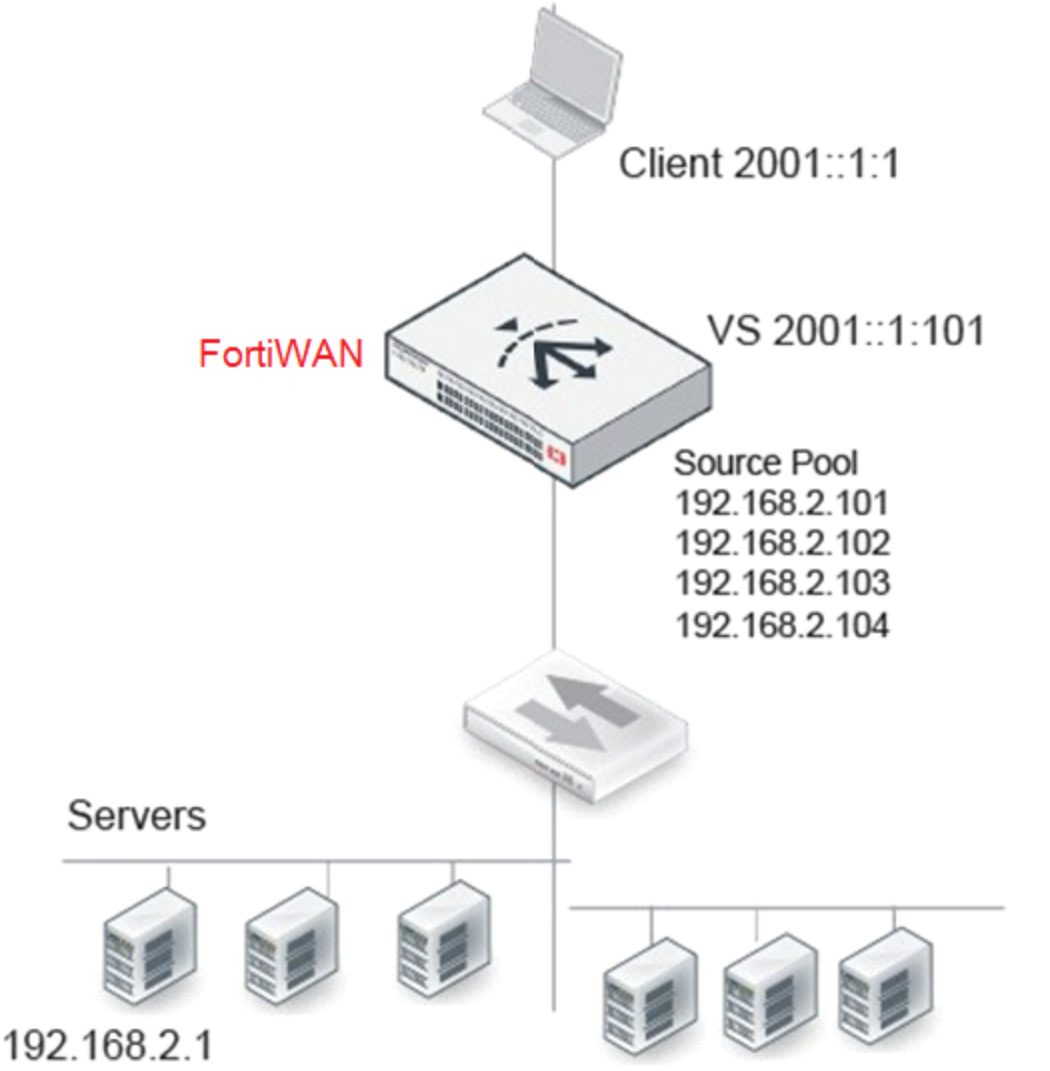

NAT64

	Ingress	Egress
SRC	2001::1:1	192.168.2.101
DST	2001::1:101	192.168.2.1

Features	Notes
Profile	Not Supported: FTP

Features	Notes
ICMP	ICMP traffic is dropped.
Security	Not Supported: IP Reputation, DoS protection, Security logs and reports

Global Load Balancing

Global load balancing (GLB) is a DNS-based solution that enables you to deploy redundant resources around the globe that you can leverage to keep your business online when a local area deployment experiences unexpected spikes or downtime. The FortiWAN system implements a hardened BIND 9 DNS server that can be deployed as the authoritative name server for the DNS zones that you configure. Zone resource records can begenerated dynamically based on the global load balancing framework. The DNS response to a client request is an ordered lists of answers that includes all available virtual servers. A client that receives DNS response with a list of answers tries the first and only proceeds to the next answers if the first answer is unreachable. The response list is based on the following priorities:

1. **Virtual server health** - Availability is determined by real-time connectivity checking. When the DNS server receives a client request, it checks connectivity for all possible matches and excludes unavailable servers from the response list.
2. **Persistence** - You can enable persistence for applications that have transactions across multiple hosts. A match to the persistence table has priority over proximity algorithms.
3. **Geographic proximity** - Proximity is determined by matching the source IP address to either the FortiGuard Geo IP database or the FortiWAN predefined ISP address book. Dynamic proximity
4. **Dynamic proximity** - Proximity is determined by application response time (RTT probes) or least connections.
5. **Weighted round robin** - If proximity algorithms are not configured or not applicable, available virtual servers are listed in order based on a simple load balancing algorithm.

The following example shows global load balancing deployment with redundant resources at data centers in China and the United States.

Figure 16 - Example Global load balancing deployment

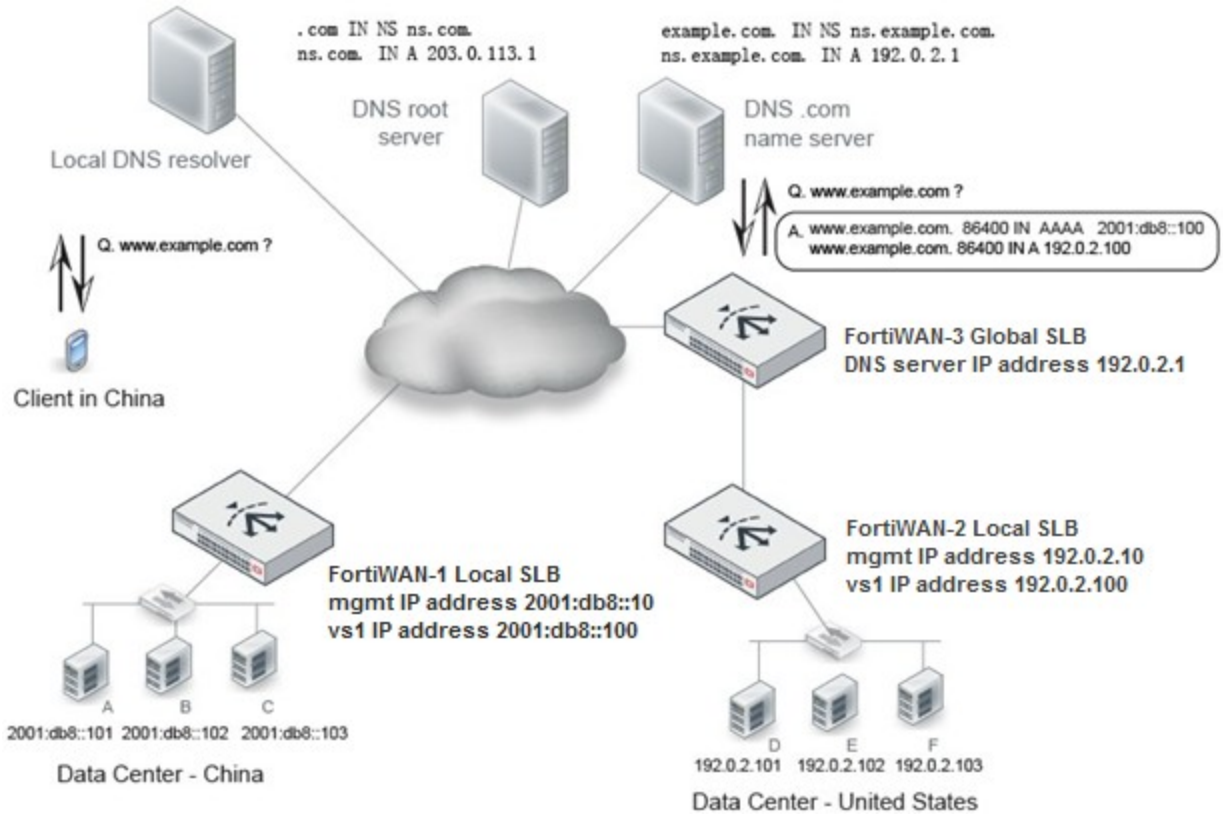

FortiWAN-1 is the local SLB for the data center in China. FortiWAN-2 is the local SLB for the data center in theUnited States. FortiWAN-3 is a global SLB. It hosts the DNS server that's authoritative for www.example.com. When a client clicks a link to www.example.com, the local host DNS resolver commences a DNS query that's ultimately resolved by the authoritative DNS server on FortiWAN-3. The set of possible answers includes the virtual servers on FortiWAN-1 or FortiWAN-2. The global load balancing framework uses health status and proximity algorithms to determine the set of answers that are returned, and the order of the answer list. Forexample, you can use the global SLB framework geoproximity feature to direct clients located in China to the virtual server in China; or if the virtual server in China is unavailable, then to the redundant resources in theUnited States.

You configure the global load balancing framework and DNS settings only on the global FortiWAN (FortiWAN-3 in the example above). The virtual server IP addresses and ports can be discovered by the FortiWAN global SLB from the FortiWAN local SLBs. The GLB DNS server uses the discovered IP addresses in the DNS response. The framework also supports third-party IP addresses and health checks for them.

The DNS server supports the following security features:

- **DNSSEC** - Domain Name System Security Extensions. DNSSEC provides authentication by associating cryptographically generated digital signatures with DNS resource record (RR) sets. The FortiWAN system makes it easy to manage the keys that must be provided to DNSparent domains and the keysthat must be imported from DNS childdomains.
- **Response rate limit** - Helps mitigate DNS denial-of-service attacks by reducing the rate at which the authoritative name servers respond to high volumes of malicious queries.
- **DNS forwarding** - In a typical enterprise local area network, the client configuration has the IP address of an internal authoritative DNS server so that requests for internal resources can be answered directly from its zone data.
- Requests for remote resources are sent to another DNS server known as a forwarder. The internal server caches the results it learns from the forwarder, which optimizes subsequent look-ups. Using forwarders reduces the number of DNS servers that must be able to communicate with Internet DNS servers.

BIND 9 reference manuals: http://www.bind9.net/manuals
RFC 1035 (DNS): http://tools.ietf.org/html/rfc1035
RFC 4033 (DNSSEC): http://tools.ietf.org/html/rfc4033

Data center

The data center configuration sets akey property, Location, which is used in the global load balancing algorithm that selects the FortiWAN in closest proximity to the client.

Prerequisites

- To select a user-defined ISP address book, create it before creating the data center configuration.
- Read-Write permission for Global Load Balancesettings.
- After you have created a data center configuration object, you can specify it in the global load balance servers configuration.

Access

- From the Dashboard, go to **Resources > Global DNS Server**, then select the **Data Center** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Location	Select a location from the location list.
Description	(optional) Type a description to help administrators know the purpose of the configuration.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Servers

In the context of the global server load balance configuration, servers are the local SLB (FortiWAN instances or third-party servers) that are to be load balanced. For FortiWAN instances, the GLB checks status and synchronizes configuration from the local SLB so that it can learn the set of virtual servers that are possible to include in the GLB virtual server pool.

The following shows configuration discovery. Placement in this list does not include them in the pool. You also must name them explicitly in the virtual server pool configuration.

Figure 17 - Virtual server discovery

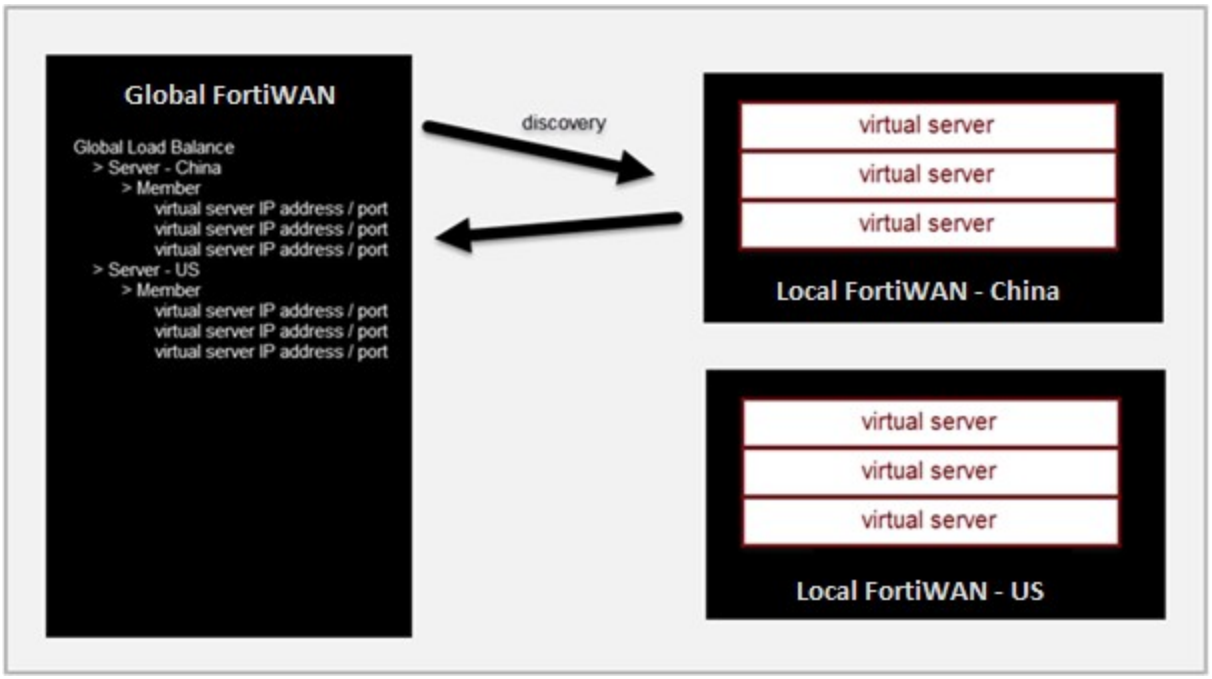

Prerequisites

- Create the data center configuration objects associated with the local SLB. See About server load balancing (SLB) on page 245.
- Create virtual server configurations on the local FortiWAN SLB. In this procedure, the global SLB discovers them. See About server load balancing (SLB) on page 245.
- Create link configuration objects on the local FortiWAN SLB if you want to configure a link health check. In this procedure, the global SLB discovers them. See Global Load Balancing on page 144.
- Read-Write permission for Global Load Balance settings.
- After creating a server configuration object, you can specify it the global load balancing virtual server pool configuration.

Access

- From the Dashboard, go to **Resources > Global DNS Server**, then select the **Server** tab.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Type	• **FortiWAN VS** - A FortiWAN instance. You must configure the IP address, as it is used for synchronization and status checks. If the management interface is unreachable, the virtual servers for that FortiWAN are excluded from DNS answers. • **Generic Host** - A third party application, controller, or server.
Data Center	Select a data center configuration object. The data center configuration object properties are used to establish the proximity of the servers and the client requests.
FortiWAN VS	
Synchronization	Toggle **ON** to synchronize the virtual server status with the local FortiWAN SLB. When enabled, synchronization occurs each time there is a change in virtual server status. Enabled by default.
Address Type	• IPv4 • IPv6
IP Address	Enter the IP address for the FortiWAN management interface. This IP address is used for synchronization and also status checks. If the management interface is unreachable, the virtual servers for that FortiWAN are excluded from DNS answers.
Generic Host	
Health Check Control	If type is Generic Host, toggle **ON** to enable health checks for the virtual server list.The health check settings at this configuration level are the parent configuration. When you configure the list, you can specify whether to inherit or override the parent configuration.

Setting	Guidelines
	This option is available only when **Generic Host** is selected. See Type above. Health checking is built-in, and you can optionally configure a gateway health check.
Health Check Relationship	• **AND** - All of the specified health checks must pass for the server to be considered available. • **OR** - One of the specified health checks must pass for the server to be considered available.
Health Check List	The items in the left pane are used for the Health Check.

- **Save**, then double-click the item from the list, or click the **Pencil** icon at the end of the row to continue entering your settings.
- Click **Add** to add a VS member.

- Click **Discover** to find members from the local FortiWAN SLB. If no window appears, there are no entries.

  After the list is populated, you can edit the configuration to add a gateway health check.
- Check **Override** to update the discovered virtual server configuration with the latest configuration information whenever you use **Discover** (for example, additions or changes to previously discovered configurations).

  Don't check this option if you want to preserve the previously discovered configuration and not have it overwritten by **Discover**.

Setting	Guidelines
Member	
Name	Enter a name that matches the virtual server configuration name on the local FortiWAN.
Address Type	• IPv4 • IPv6
IP Address	Enter the IP address of the virtual server. GLB answers this IP address for DNS requests if this virtual server is selected from the virtual server pool associated to the queried domain.

Setting	Guidelines
	By default, GLB answers this IP address for DNS requests coming from both internal and external networks. When Custom External IP is set, this IP will be answered only for DNS requests coming from internal networks.
Custom External	Toggle ON to have the GLB answer this IP address for DNS requests coming from external IP networks.
Port	Virtual server port. The default if 80. The valid range is 1 to 65535.
Protocol	• TCP • UDP The default is TCP.
WAN Link	The outbound WAN link of the virtual server. Usually, this is the WAN link that connects to the same interface that the virtual server is listening on. Further, you can associate this with a link object to establish the proximity of the servers and the client requests. See About global load balancing on page 248. This is used only when type is set as FortiWAN VS, not for Generic Host.
Health Check Inherit	Toggle **ON** to inherit the health check settings from the parent configuration, if the type is Generic Host.
Health Check Control	This option is available only when Health Check Inherit is disabled. Toggle **ON** to enable this option.
Health Check Relationship	• **AND** - All of the specified health checks must pass for the server to be considered available. • **OR** - One of the specified health checks must pass for the server to be considered available.
Health Check List	The items in the left pane are used for the Health Check.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.

- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** ✖ icon at the end of the row.

Link (GLB)

On this page you can set key properties: ISP and ISP province, which are used in global load balancing to select the closest FortiWAN to the client.

Configure this to associate a ISP profile with the WAN link of a virtual server for the proximity of the servers and the client requests.

Prerequisites

- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Resources > Global DNS Server**, then select the **Link** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Configuration name. Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Data Center	Select a data center from the list. You must first configure the data centers. See Data center on page 146.
ISP	Select an ISP from the list. See Resources > Address > ISP Address.
ISP Province	Select an ISP province from the list. See Resources > Address > ISP Address.
Server Link Mapping	
Server	Select a server. See Servers on page 147.
WAN Link	The WAN link of FortiWAN that can access to the virtual server.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Virtual server pool

The virtual server pool configuration defines the set of virtual servers that can be matched in DNS resource records, so it should include, for example, all the virtual servers that can be answers for DNS requests to resolve www.example.com.

You also specify the key parameters of the global load balancing algorithm, including proximity options, status checking options, load balancing method, and weight.

The DNS response is an ordered list of answers. Virtual servers that are unavailable are excluded. Available virtual servers are ordered based on the following priorities:

1. Geographic proximity
2. Dynamic proximity
3. Weighted round robin

A client that receives DNS response with a list of answers tries the first and only proceeds to the next answers if the first answer is unreachable.

Prerequisites

- Configure the GLB Servers. See DNS Server (DNS zones) on page 180.
- Read-Write permission for Global Load Balance settings.

After you have created a virtual server pool configuration object, you can specify it in the global load balancing host configuration.

Access

- From the Dashboard, go to **Resources > Global DNS Server**, then select the **Virtual Server Pool** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. Reference this name in the host configuration.
Persistence	Toggle **ON** to enable the persistence table. Disabled by default. If you enable persistence, the client source address is recorded in the persistence table, and subsequent requests from the same network or the same host or domain are sent an answer with the virtual servers listed in the same order (unless a server becomes unavailable and is therefore omitted from the answer).
Respond Single Record	Toggle **ON** to send a single record in response to a query. Disabled by default. By default, the response is an ordered list of records.
Preferred	Choose from the drop-down list. • **WRR** - Weighted Round Robin. • **Geo** - If selected, virtual servers with the same GEO information as the local DNS address will be responded. • **Geo-ISP** - If selected, virtual servers with the same ISP information as the local DNS address will be responded first, and virtual servers with the same GEO information as the local DNS address will be responded second. • **RTT** - Virtual servers with the shortest latency link or closest to the data center will be responded. • **VS-Least-Connnections** - Virtual servers with the least connections will be responded. • **VS-Connection-Limit** - Virtual servers with higher connection limit setting will be responded. • **VS-Traffic** - Virtual servers with the lowest traffic will be responded. • **Link-Upstream-Traffic** - The virtual server with the least

Setting	Guidelines
	upstream traffic on the associated WAN link will be responded. • **Link-Downstream-Traffic** - The virtual server with the least downstream traffic on the associated WAN link will be responded. • **Link-Total-Traffic** - The virtual server with the least total traffic on the associated WAN link will be responded.
Check Server Status	Toggle **ON** to enable polling of the local FortiWAN SLB. If the server is unresponsive, its virtual servers are not selected for DNS answers.
Check Virtual Server Existence	Toggle **ON** to check whether the status of the virtual servers in the virtual server list is known. Virtual servers with unknown status are not selected for DNS answers.
Member	
Server	Select a GLB Server configuration object.
Server Member	Select the name of the virtual server that's in the servers virtual server list configuration.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Backup	Enable to designate the member as a backup. Backup members are inactive until all main members are down.

- **Save**, then click **Add** to continue entering Member data.

Setting	Guidelines
Server	Select a GLB Server configuration object.
Server Member	Select the name of the virtual server that's in the servers virtual server list configuration.
TTL	**Time to Live** - Enter the number of seconds to keep the packet.

Setting	Guidelines
	The default is -1. The valid range is -1 to 2147483647. -1 means it uses the zone level TTL.
Weight	Assigns relative preference among members—higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Backup	Toggle **ON** to enable to designate the member as a backup. Backup members are inactive until all main members are down.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Remote DNS server - Resources

The remote server configuration is used to create a list of DNS forwarders. DNS forwarders are commonly used when you don't want the local DNS server to connect to Internet DNS servers. For example, if the local DNS server is behind a firewall and you don't want to allow DNS through that firewall, you implement DNS forwarding to a remote server that's deployed in a DMZ or similar network region that can contact Internet DNS servers.

Prerequisites

- Good understanding of DNS and knowledge of the remote DNS servers that are used to communicate with Internet domain servers.
- Read-Write permission for Global Load Balance settings. See DNS Server (DNS zones) on page 180.

After you have configured remote DNS servers, you can select them in DNS zone and DNS policy configurations.

Access

- From the Dashboard, go to **Resources > Global DNS Server**, then select the **Remote DNS Server** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Address Type	IPv4 IPv6
IP	IP address of the remote DNS server.
Port	Port number the remote server uses for DNS. The default is 53.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Services

FortiWAN enables you to configure link load balance, global DNS servers, and virtual servers.

Link load balance (LLB) - Services

Link load balancing (LLB) manages traffic over multiple internet service providers (ISP) or wide area networks (WAN). If your ISP uses billing tiers based on bandwidth rate or peak/off-peak hours, you can provision multiple links, reducing the risk of outages, obtaining additional bandwidth for peak events, potentially saving costs.

In most cases, configure link load balancing for outgoing traffic. Outbound traffic might be user or server traffic that's routed from your local network through your ISP transit links, leased lines, or other WAN links to destinations on the Internet or WAN. Configure link policies that select the gateway for outbound traffic.

When the FortiWAN system receives outbound traffic that matches a source-destination-application tuple that you configure, it forwards it to an outbound gateway link according to system logic and policy rules that you specify.

LLB supports load balancing among Underlay link groups or Overlay link (IPsec tunnels) groups.

Flow Policy

Configure your flow policy from this page.

A flow policy matches traffic to rules that select a underlay link group or a overlay link group and underlay default link group.

- The policy uses a matching tuple: source, destination, application, and schedule. The policy match is a Boolean AND—All must match for the rule to be applied.
- The elements of the tuple support specification by group objects. This is a Boolean OR—If source IP address belongs to member 1 OR member 2, then source matches.
- The logical combinations enable you to subscribe multiple address spaces or applications to a group of links, and create load balancing rules on that group basis.

Prerequisites

- Configure any address, application, and schedule objects that you want to use as match criteria for your policy.
- Configure a link group.
- Read-Write permission for Services module settings.

Access

- From the Dashboard, go to **Services > Link Load Balance**, then select the **Flow Policy (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.
- To use the predefined configuration, from **Default Link Group**, choose `underlay_all_wrr`.
- To create a new policy, click **Add** to open the configuration editor.

Settings	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Ingress interface	This field only appears then Ingress Type **Interface** is chosen. Choose a port from the drop down list.
Ingress Type	• **Interface** - The match traffic from an interface. • **Zone** - The match traffic from a source zone.
Source Type	• **Address** - see Address on page 103. • **Address Group** - see Address group on page 104.
Source Address Source Address Group	Choose from the drop down list.
Source Zone	This field only appears when Source Type **Address Group** is chosen. • WAN • LAN • DMZ
Destination Type	• **Address** - see Address on page 103. • **Address Group** - see Address group on page 104.
Destination Address	Choose from the drop down list.

Settings	Guidelines
Destination Address Group	
Service Type	• **App** - see Applications on page 107. • **App Group** - see Application groups on page 110.
Application Application Group	Choose from the drop down list.
Schedule	Choose from the drop down list. See Schedules and schedule groups on page 111.
Link Group	Choose from the drop down list. Don't use the default link group. See Link group on page 116.
Tunnel Quality Policy	
Tunnel Quality Policy	• **Disable** - Select to turn off the tunnel quality policy. • **Enable** - Select to use the tunnel quality policy.
Tunnel Quality Threshold Rtt	Enter the threshold number. The default is 150. The default is 150 Milliseconds. The valid range is 1 to 1000.
Tunnel Quality Threshold Jitter	Enter the threshold jitter number. The default is 150. The default is 150 Milliseconds. The valid range is 1 to 1000.
Tunnel Quality Threshold Pack Loss	Enter the number of lost packets to tolerate. The default is 5. The valid range is from 1 to 99.
Traffic Log	Toggle **ON** to log any traffic matching the link flow policy.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Source NAT

Use source NAT (SNAT) when clients have IP addresses from private networks. This ensures you don't have multiple sessions from different clients with source IP 192.168.1.1, for example. Or, you can map all client traffic to a single source IP address because a source address from a private network isn't meaningful to the FortiWAN system or backend servers.

Shown below, the SNAT rule matches the source and destination IP addresses in incoming traffic to the ranges specified in the policy. If the client request matches, the system translates the source IP address to an address from the SNAT pool. In this example, a client with private address 192.168.1.1 requests a resource from the virtual server address at 192.0.2.1 (not the real server address 10.0.0.1; the real server address isn't published). The two rule conditions match, so the system translates the source IP to the next address in the SNAT pool—10.1.0.1. SNAT rules don't affect destination addresses, so the destination address in the request packet is preserved.

The system maintains this NAT table and performs the inverse translation when it receives the server-to- client traffic. Be sure to configure the backend servers to use the FortiWAN address as the default gateway so that server responses are also rewritten by the NAT module.

Note: This SNAT feature isn't supported for traffic to virtual servers. Use the virtual server SNAT feature instead.

Figure 18 - SNAT

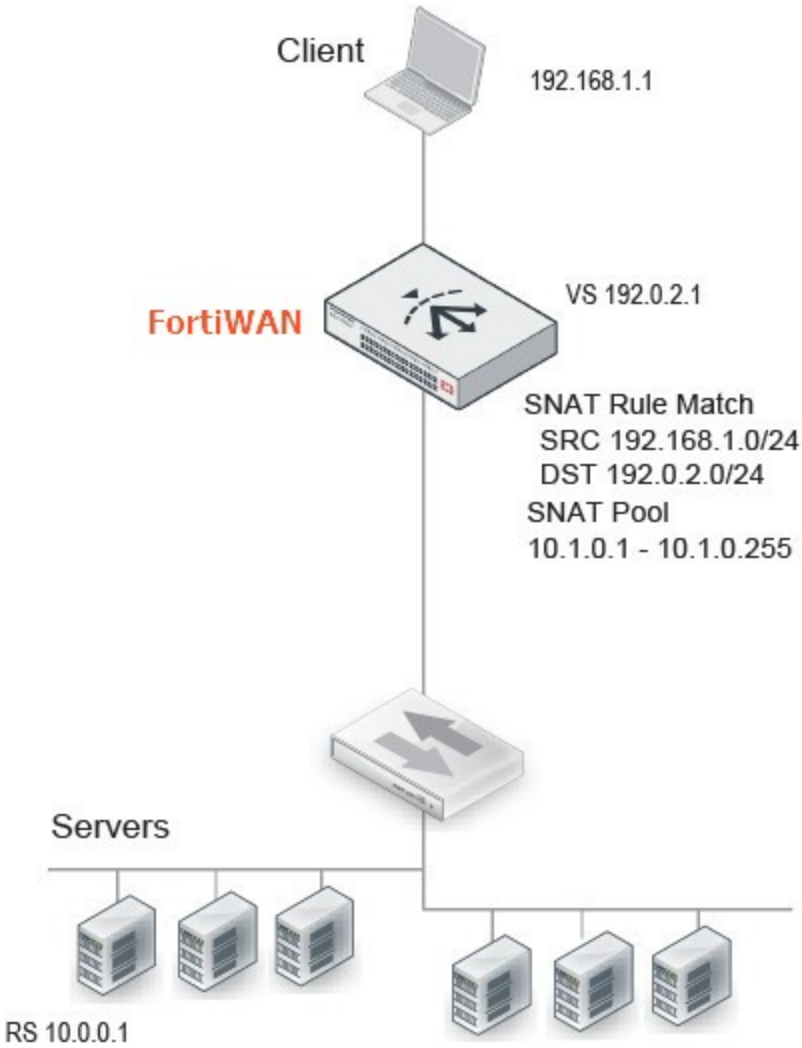

SNAT		Ingress	Egress
	SRC	192.168.1.1	10.1.0.1
	DST	192.0.2.1	192.0.2.1

Prerequisites

- Know the IP addresses your organization has provisioned foryour NAT design.
- Read-Write permission for Networking settings.

Access

- From the Dashboard, go to **Services > Link Load Balance**, then select the **Source NAT** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Source	Choose the Address/mask notation to match the source IP address in the packet header.
Destination	Choose the Address/mask notation to match the destination IP address in the packet header.
Transiation Type	• IP Address • Use Interface IP Address • No Nat - FortiWAN will not translate the source address
Pool Address Range	This option applies only when the Translation Type is set to IP address. Enter the first IP address in the SNAT pool.
To	This option applies only when the Translation Type is set to IP address. Enter the last IP address in the SNAT pool.
From Type	• Address • Address Group
From Group	This option applies only when the From Type is set to Address Group. Select a address group as the source addresses.
To Type	• Address • Address Group
To Group	This option applies only when the To Type is set to Address Group. Select a address group as the destination addresses.
Application Type	• Application • Application Group
Application	Select a application which you want to make SNAT.
Application Group	This option applies only when the Application is set to Application Group. Select a application group which you want to make SNAT.

Setting	Guidelines
Schedule	Select the schedule object.
Out Type	• Interface • Link
Out Interface	Choose an out interface from the drop down list.
Out Link	This option applies only when the Out type is set to Out Link. Choose a link from the drop down list.
Trans to Random Port	Toggle **ON**. FortiWAN translates the packets with a random source port.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

1-to-1 NAT

Use 1-to-1 NAT when you want to publish public or “external” IP addresses for FortiWAN resources, but want the communication among servers on the internal network to be on a private or “internal” IP address range.

The NAT configuration example shown below assigns both external and internal (or “mapped”) IP addresses to Interface 1. Traffic from the external side of the connection (such as client traffic) uses the external IP address and port. Traffic on the internal side (such as the virtual server communication with real servers) uses the mapped IP address and port.

1-to-1 NAT is supported for traffic to virtual servers. The address translation occurs before the FortiWAN has processed its rules, so FortiWAN server load balancing policies that match source address (such as content routing and content rewriting rules) should be based on the mapped address space.

The system maintains this NAT table and performs the inverse mapping when it sends traffic from the internal side to the external side.

Figure 19 - 1-to-1 NAT

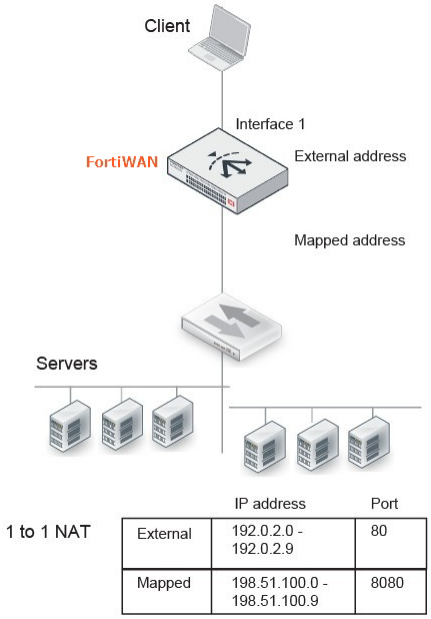

Prerequisites

- Know the IP addresses your organization has provisioned for your NAT design.
- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Services > Link Load Balance**, then select the **1-to-1 NAT** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
External Interface	Choose the Interface that receives traffic from the drop down list.
External Address	Enter the first address in the range. The last address is calculated after you enter the mapped address range.
Mapped Address Range	Enter the first and last addresses in the range.
Port Forwarding	Toggle **ON** to enable.
Protocol	•ICMP •TCP •UDP
External Port	Enter the first port number in the range. The last port number is calculated after you enter the mapped port range.
Mapped Port Range	Enter the first and last port numbers in the range.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Bandwidth

A Bandwidth policy assigns traffic to the traffic shaper.

The FortiWAN system does not provision bandwidth based on the TOS bits (also called differentiated services) in the IP header to control packet queuing. Instead, the system provisions bandwidth based on a source-destination-application matching tuple that you specify.

The bandwidth policy feature isn't supported for traffic to virtual servers.

Workflow

1. Configure a traffic shaper for bandwidth management. See Traffic Shaper on page 122.
2. Configure a bandwidth policy.

Prerequisites

- Have a good understanding of traffic in your network that requires bandwidth provisioning.
- Create the address configuration objects and application configuration objects that define the matching tuple for bandwidth policy.
- Create a traffic shaper configuration object.
- Read-Write permission for Link Load Balance settings.

Access

- From the Dashboard, go to **Services > Link Load Balance**, then select the **Bandwidth** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Status	Toggle **ON** to enable the bandwidth policy.
Traffic Shaper	Choose the traffic shaper that will be used for packets that match the bandwidth policy criteria. This affects uploads or outbound traffic.
Reverse Traffic Shaper	Choose the traffic shaper that will be used for packets that match the bandwidth policy criteria. This affects downloads or inbound traffic.
Application Type	
	• **Application** - Choose an application object to use to form the matching tuple. • **Application Group** - Choose an application group object to use to form the matching tuple.
Source Type	
	• **Address** - Choose a source address object to use to form the matching tuple. • **Address Group** - Choose a source address group object to use to form the matching tuple.
Destination Type	
	• **Address** - Choose a destination address object to use to form the matching tuple. • **Address Group** - Choose a destination address group object to use to form the matching tuple.
In type	
	• **Interface** - the interface that receives traffic.

Setting	Guidelines
	• **Zone** - the zone that receives traffic. • **Link** - the link that receives traffic.
Out type	
	• **Interface** - the interface that forwards traffic. • **Zone** - the zone that forwards traffic. • **Link** - the link that forwards traffic.
Schedule	Choose the schedule.
Drop Log	Toggle ON to enable the drop log. When enabled, if the packet is dropped due to bandwidth limit, the log is triggered.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Connection limit

The connection limit policy allows or denies traffic based on a matching tuple: source address, destination address, and service; and connection count. The purpose is to detect anomalous connection requests.

The limit you specify can be based on the following counts:

- Count of concurrent sessions that match the tuple.
- Count of concurrent sessions from a single host that match the tuple.

The FortiWAN system evaluates connection limit policy rules before other rules. It matches traffic against the connection limit table, beginning with the first rule. If no rule matches, the connection is forwarded for further processing. If a rule matches, and the limit has not been reached, the connection is forwarded for further processing. If a rule matches and the limit has been reached, the connection is dropped.

By default, if connection limit rules are not configured, the system does not perform connection limit policy processing.

The purpose of the connection limit is different from the virtual server connection limit. The connection limit setting is a security setting; the virtual server connection limit is a capacity setting.

Prerequisites

- Know the capacity of your backend servers.
- Create the address configuration objects and service configuration objects that define the matching tuple in your connection limit rules.

Access

- From the Dashboard, go to **Services > Link Load Balance**, then select the **Connection Limit** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
Status	Toggle ON to monitor bandwidth status.
Traffic Shaper	Select the traffic shaper that will be used for packets that match the bandwidth policy criteria. This affects uploads or outbound traffic. See Traffic Shaper on page 122.
Reverse Traffic Shaper	Select the traffic shaper that will be used for packets that match the bandwidth policy criteria. This affects downloads or inbound traffic.
Application Type	• **Application** - Select to use an application to form the matching tuple. See Address on page 103. • **Application Group** - Select to use an application group to form the matching tuple. See Address group on page 104.
Application Application Group	Choose an application or application group which you want to make SNAT.
Source Address	Select a source address to use to form the matching tuple.
Destination Address	Select a destination address to use to form the matching tuple.
In Interface	Select the interface that receives traffic. See Interface settings on page 63.
Out Interface	Select the interface that forwards traffic.
In type	Select the interface type for inbound traffic. • **Interface** - Select the interface that receives traffic. • **Zone** - This option applies only when the In Type is set to Zone. Select the zone that receives traffic. • **Link** - This option applies only when the In Type is set to Link. Select the link that receives traffic.

Setting	Guidelines
Out type	Select the interface type for outbound traffic. • **Interface** - Select the interface that forwards traffic. • **Zone** - This option applies only when the Out Type is set to Zone. Select the zone that forwards traffic. • **Link** - This option applies only when the Out Type is set to Link. Select the link that forwards traffic.
In Zone Out Zone	Select a zone. • WAN • LAN • DMZ
In Link Out Lint	Choose a link from the drop down list. See Links - underlay or overlay on page 76.
Source Type	• **Address**- Address/mask notation to match the source IP address in the packet header. For example, 192.0.2.0/24. • **Address Group** - This option applies only when the Source Type is set to Address Group. Select a address group as the source addresses
Destination Type	• **Address** - Address/mask notation to match the destination IP address in the packet header. For example, 10.0.2.0/24. • **Address Group** - This option applies only when the Destination Type is set to Address Group. Select a address group as the destination addresses
Schedule	Choose from the drop-down list. See Schedules and schedule groups on page 111.
Drop Log	Toggle ON to enable the drop log. When enabled, if the packet is dropped due to connection limit, the log is triggered.

Firewall

A firewall policy is a filter that allows or denies traffic based on a matching tuple: source address, destination address, and service. By default, firewall policy rules are stateful: if client-to-server traffic is allowed, the session is maintained in a state table, and the response traffic is allowed.

The FortiWAN system evaluates firewall policies before other rules. It matches traffic against the firewall policy table, beginning with the first rule. If a rule matches, the specified action is taken. If the session is denied by a firewall policy rule, it is dropped. If the session is accepted, system processing continues.

By default, if firewall rules are not configured, the system does not perform firewall processing; all traffic is processed as if the system were a router, and traffic is forwarded according to routing and other system rules.

You don't need to create firewall rules for routine management traffic associated with the management port or HA ports. The interface “allow access” option enables permitted protocols. The system automatically permits from-self traffic, such as health check traffic, and expected responses.

Prerequisites

- Have a good understanding and knowledge of firewalls.
- Create the address configuration objects and service configuration objects that define the matching tuple in your firewall policy rules.

Access

- From the Dashboard, go to **Services > Firewall**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Default Action	Action when no rule matches or no rules are configured: • **Deny** - Drop the traffic. • **Accept** - Allow the traffic to pass the firewall.
Rule	
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name.
In Interface	
Out Interface	
Source Address	
Destination Address	
Action	• **Deny** - Drop the traffic. • **Accept** - Allow the traffic to pass the firewall.
Source Type	• Address • Address Group
Source Address	Address/mask notation to match the source IP address in the packet header. For example, 192.0.2.0/24.
Source Address Group	This option applies only when the Source Type is set to Address Group. Select a address group as the source address.
Destination Type	• Address • Address Group
Destination Address	Address/mask notation to match the destination IP address in the packet header. For example, 10.0.2.0/24.
Destination Address Group	This option applies only when the Destination Type is set to Address Group. Select a address group as the destination address.

Setting	Guidelines
In type	Select the interface type for inbound traffic. • Interface • Zone • Link
In Interface	Select the interface that receives traffic.
In Zone	This option applies only when the In Type is set to Zone. Select the zone that receives traffic.
In Link	This option applies only when the In Type is set to Link. Select the link that receives traffic.
Out type	Select the interface type for outbound traffic. • Interface • Zone
Out Interface	Select the interface that forwards traffic.
Out Zone	This option applies only when the Out Type is set to Zone. Select the zone that forwards traffic.
Setting	Guidelines
Application Type	• Application • Application Group
Application	Select a application which you want to make SNAT
Application Group	This option applies only when the Application is set to Application Group. Select a application group which you want to make SNAT
Schedule	Select the schedule from the drop down list.
Deny Log	Toggle **ON** to enable the deny log. The log records each time traffic is dropped.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.

- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** ✖ icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.
- Use the arrow buttons 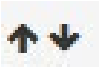 to change the order of the rules. Rules are evaluated from top to bottom. The first rule that matches is applied and subsequent rules are not evaluated.

DNS Server (DNS zones)

The DNS zone configuration is the key to the global load balancing solution. This configuration contains the key DNS server settings, including:

- Domain name and name server details.
- **Type** - Whether the server is the master or a forwarder.
- **DNSSEC** - Whether to use DNSSEC.
- **DNSRR** records - The zone configuration contains resource records (RR) used to resolve DNS queries delegated to the domain by the parent zone.

Enter different DNS server settings for each zone you create. For example, the DNS server can be a master for one zone and a forwarder for another zone.

The general settings configuration specifies the interfaces that listen for DNS requests. By default, the system listens on the IPv4 and IPv6 addresses of all configured interfaces for DNS requests.

The other settings in the general settings configuration are applied when traffic does not match a Global DNS policy.

Prerequisites

- A good understanding of DNS and knowledge of the DNS deployment in your network.
- Permission to create authoritative DNS zone records for your network.
- Read-Write permission for Global Load Balance settings. After you have configured a DNS zone, you can select it in the DNS policy configuration.

Access

- From the Dashboard, go to **Services > Global DNS Server**, then select the **DNS Server (default)** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. Reference the name in the global DNS policy configuration. • FortiWAN supports third-party domain names.
Server Type	• **Internal**–Listen DNS queries on all the LAN ports. • **External**–Listen DNS queries on all the WAN ports. • **All** – Listen DNS queries on all network interfaces.
Type	• **Forward** - forward resolution to an IP address. • **Reverse** - reverse resolution of an IP address.
Domain Name	Enter a domain name, which must end with a period. For example: example.com.
General Configurations	
TTL	The $TTL directive at the top of the zone file (before the SOA) gives a default TTL for every RR without a specific TTL set. The default is 86,400. The valid range is -1 to 2147483647.
Responsible Mail	(required) Enter the email address of the person responsible for this zone. Don’t use “@”; instead, use a dot, because @ has other uses in the zone file. For example, *admin.example.com.*, including the trailing dot. Email, however, is sent to admin@example.com.
Primary Server Name	(required) Sets the server name in the SOA record.
Primary Server Address (IPv4)	The IPv4 address of the primary server. For example, 192.0.2.1.
Primary Server Address (IPv6)	The IPv6 address of the primary server.
Reverse	

Setting	Guidelines
Ptr Address Type	• IPv4 • IPv6
Negative TTL	Negative Time to Live - Enter the number of seconds to keep the packet. The default is 86400. The range is from 1 to 2147483647.
Refresh	Enter the number of seconds between refresh. The default is 86400. The range is from 1 to 2147483647.
Retry	Enter the number of seconds before retry. The default is 86400. The range is from 1 to 2147483647.
Expire	Enter the number of seconds the packet will expire. The default is 86400. The range is from 1 to 2147483647.
Serial Number	Enter the serial number. The default is 2017120701. The range is 1 to 2147483647.
DNSSEC Settings	
	• Toggle **ON** to use External DNSSEC Status. • Select RSA-SHA1 to use this DNSSEC Algorithm.

- **Save**, then double-click the item from the list, or click the **Pencil** icon at the end of the row.

DNSSEC Settings	
DNSSEC	Toggle **ON** to enable DNSSEC.
DNSSEC Algorithm	Only RSA-SHA1 is supported.
Renew Private Keys	Click to renew the private keys.
Export DNSSEC Files	Click to export the DNSSEC files.
A/AAAA Record	
Hostname	(required) The hostname part of the FQDN, such as www. You can specify the @ symbol to denote the zone root. The

	value substituted for @ is the preceding $ORIGIN directive.
Type	• **Static** for standard DNS A/AAAA records with static IP addresses. • **Dynamic** for associating with GLB to answer dynamic IP addresses according to status of SLB virtual servers.
TTL	**Time to Live** - Enter the number of seconds to keep the packet. The default is -1. The valid range is -1 to 2147483647. -1 means it uses the zone level TTL.
Source Type	• **IPv4** for A records. • **IPv6** for AAAA records.
Static A/AAAA Record	
Weight	Assigns relative preference among members - higher values are more preferred and are assigned connections more frequently. The default is 1. The valid range is 1-255.
Address	Enter the IP address of the virtual server. GLB answers this IP to queries coming from both internal and external networks by default. When Custom External is enabled and set, this IP is only for queries coming form internal networks.
Custom External	Toggle ON to set an external IP.
External IP	GLB answers this IP from queries coming from external networks.
Dynamic A/AAAA Record	
Virtual Server Pool	Select a defined virtual server pool so that GLB answers the IP address of a virtual server from the pool dynamically with the specified algorithm. See Configuring data centers, Configuring servers and Configuring virtual server pools.
CNAME Record	
Alias	An alias name to another true or canonical domain name (the target). For instance, www.example.com is an alias for

	example.com.
Target	The true or canonical domain name. For instance, example.com.
TTL	**Time to Live** - Enter the number of seconds to keep the packet. The default is -1. The valid range is -1 to 2147483647. -1 means it uses the zone level TTL.
NS Record	
Domain Name	(required) The domain for which the name server has authoritative answers, such as example.com. FortiWAN supports third-party domain names.
Hostname	(required) The hostname part of the FQDN, such as ns.
TTL	**Time to Live** - Enter the number of seconds to keep the packet. The default is -1. The valid range is -1 to 2147483647. -1 means it uses the zone level TTL.
Type	• IPv4 • IPv6
Address	Enter the IP address of the name server.
MX Record	
Domain Name	(required) The domain for which the name server has authoritative answers, such as example.com. FortiWAN supports third-party domain names.
Hostname	(required) The hostname part of the FQDN.
TTL	**Time to Live** - Enter the number of seconds to keep the packet. The default is -1. The valid range is -1 to 2147483647. -1 means it uses the zone level TTL.
Priority	Preference given to this RR among others at the same owner. Lower values have greater priority.
Type	• IPv4

	• IPv6
Address	Enter the IP address.
TXT Record	
Name	Hostname. TXT records are name-value pairs that contain human readable information about a host. The most common use for TXT records is to store SPF records.
Text	Comma-separated list of name=value pairs. An example SPF record has the following form: `v=spf1 +mx a:colo.example.com/28 -all` If you complete the entry from the the Web UI, do not put the string in quotes. (If you complete the entry from the CLI, you do put the string in quotes.)
SRV Record	
Host Name	The host name part of the FQDN,such as www.
Priority	A priority assigned to the target host: the lower the value, the higher the priority.
Weight	A relative weight assigned to a record among records of the same priority: the greater the value, the more weight it carries.
Port	The TCP or UDP port on which the service is provided.
Target Name	The canonical name of the machine providing the service.
DName Record	
Alias	
Target	(required)
TTL	**Time to Live** - Enter the number of seconds to keep the packet. The default is -1. The valid range is -1 to 2147483647. -1 means it uses the zone level TTL.

- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** ✖ icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

DNS Settings (dynamic proximity) - Services

Use this page to configure dynamic proximity. Dynamic proximity is used to order DNS lookup results based on round-trip time (RTT) for ICMP or TCP probes sent by the local SLB to the DNS resolver that sent the DNS request.

The system caches the RTT results for the period specified by the timeout. When there are subsequent requests from clients that have a source IP address within the specified netmask, the RTT is taken from the results table instead of a new, real-time probe. This reduces DNS response time.

Prerequisites

- Read-Write permission for Global Load Balance settings.

The settings you configure are applied if dynamic-proximity is enabled in the virtual server pool configuration.

Access

- From the Dashboard, go to **Services > Global DNS Server**, then select the **DNS Settings** tab.

Settings

- Set your options as needed, then **Save**.

Setting	Guidelines
Internal DNS Server	Toggle **ON** to enable an internal DNS server.
External DNS Server	Toggle **ON** to enable an external DNS server.
Internal Request Forward	Toggle **ON** to forward internal requests.
External Request Forward	Toggle **ON** to forward external requests.
Internal Response Rate Limit	Enter the rate limit. The default is 0. The valid range is 1 to 102400.
External Response Rate Limit	Enter the rate limit. The default is 0. The valid range is 1 to 102400.
Traffic log	Toggle **ON** to enable the traffic log. See Traffic log on page 211.
Listen on all interface	Toggle ON to enable listening on all network interfaces.
Persistence	
Proximity Mask Length	Number of IPv4 netmask bits that define network affinity for the RTT table. The default is 24. The valid range is 1 to 32. For example, if the GLB records an RTT for a client with source IP address 192.168.1.100, the record is stored and applies to all requests from the 192.168.1.0/24 network.
Proximity Mask Length6	Number of IPv6 netmask bits that define network affinity for the RTT table. The default is 64. The valid range is 1 to 128.
Proximity Aging Period	RTT results are cached. This setting specifies the length of time in seconds for which the RTT cache entry is valid. The default is 86400. The valid range is 60 to 2,592,000 seconds.
Proximity Settings	
Proximity Proto	• **ICMP**- Use ICMP to detect routes. Calculate proximity by the

Setting	Guidelines
	smaller RTT. • **ICMP and TCP**- Some hosts do not respond to ICMP requests. Specify this option to use both ICMP and TCP to detect routes and RTT. For TCP detection, a SYN packet is sent to port 53. A connection refused or connection reset by the destination is treated as successful detection
Proximity Retry Num	The retry count if the probe fails. The default is 3. The valid range is 1 to 10 times.
Proximity Retry Internal	Interval between retries if the probe fails. The default is 3. The valid range is 1 to 3600 seconds.
Proximity Mask Length	Number of IPv4 netmask bits that define network affinity for the RTT table. The default is 24. The valid range is 1 to 32. For example, if the GLB records an RTT for a client with source IP address 192.168.1.100, the record is stored and applies to all requests from the 192.168.1.0/24 network.
Proximity Mask Length6	Number of IPv6 netmask bits that define network affinity for the RTT table. The default is 64. The valid range is 1 to 128.
Proximity Aging Period	RTT results are cached. This setting specifies the length of time in seconds for which the RTT cache entry is valid. The default is 86400. The valid range is 60 to 2,592,000 seconds.

Virtual server settings

This topic describes how to configure a virtual server.

The virtual server configuration mainly supports Layer 4 application delivery control:

- **Layer 4** - Persistence, load balancing, and network address translation are based on Layer 4 objects, such as source and destination IP address.

Prerequisites

- Have a deep understanding of the backend servers and your load balancing objectives.
- Configure a real server pool (required) and other configuration objects that you can incorporate into the virtual server configuration, such as persistence rules, user-defined profiles and NAT pools if you are deploying NAT. See Real server pool on page 133.
- Read-Write permission for Load Balance settings.

Unlike virtual IPs on FortiGate or virtual servers on FortiWeb, virtual servers on FortiWAN are activated as soon as you have configured them and set their status to Enable, not by selecting them in a policy.

After the configuration resources for virtual server are prepared (see Virtual Servers on page 126), you can start setting the virtual server on the FortiWAN interface to receive service requests. Before starting to configure the virtual server, know that:

- The FortiWAN virtual server supports only standard port 21 FTP for both active and passive modes; other ALG applications are not supported.
- To implement FTP passive mode, the FortiWAN virtual server listens on any port higher than 1024 and take the received packets (with destination port higher than 1024) as connectivity of a data link. This might break in on other services (such as SNAT and FortiWAN GLB management) that run via the same IP address as the virtual server. You might assign an exclusive virtual IP address for the virtual server to make sure no other services share the same IP address.

Access

- From the Dashboard, go to **Services > Virtual Server**.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter a unique name. Valid characters are A-Z, a-z, 0-9, _, and -. No spaces. Once you **Save**, you can't change the name. This name appears in reports and in logs as the SLB “policy”.
Interface	Network interface that receives client traffic for this virtual server.
Address Type	• IPv4 • IPv6 - IPv6 is not supported for FTP profile.
Address	Enter the IP address provisioned for the virtual server when IP Type is set to **Specify**. You can specify a IP address which is not physically bound on the interface via **Networking > Interface**, FortiWAN takes the specified IP address as virtual IP on the interface; you just make sure clients can access to this IP.
Packet Forwarding Method	For Layer 4 virtual servers, select one of the following packet forwarding methods: • **Direct Routing** - Forwards the source and destination IP addresses with no changes. **Note**: For FTP profiles, when Direct Routing is selected, you must also configure a persistence method. • **NAT** - Replaces the destination IP address with the IP address of the backend server selected by the load balancer. The destination IP address of the initial request is the IP address of the virtual server. Be sure to configure FortiWAN as the default gateway on the backend server so that the reply goes through FortiWAN and can also be translated. • **Full NAT** - Replaces both the destination and source IP addresses. IPv4 to IPv4 or IPv6 to IPv6 translation. • **NAT46** - Replaces both the destination and source IP addresses, translating IPv4 addresses to IPv6 addresses. • **NAT64** - Replaces both the destination and source IP addresses, translating IPv6 addresses to IPv4 addresses. For Full NAT, NAT46, and NAT64, the source IP address is

Setting	Guidelines
	replaced by an IP address from the pool you specify. The destination IP address is replaced with the IP address of the backend server selected by the load balancer
Port	Select a port number to listen for client requests. If a Layer 2 virtual server is assigned a network interface that uses port 80 or 443, ensure that the HTTPS and HTTP administrative access options are not enabled for the interface.
Load Balance Profile	Select a predefined or user-defined profile configuration object. See Link load balance (LLB) - Resources on page 113.
Server Pool Name	Select a real server pool configuration object. See Real server pool on page 133.
Traffic Log	Enable to record traffic logs for this virtual server.
Comments	A string to describe the purpose of the configuration, to help you and other administrators more easily identify its use.
Algorithm	Select the algorithm to distribute service requests to the appropriate real server from the real server pool. The available algorithms are: • **Round Robin**: Weighted round robin to select a real server from the pool for a request. Weight of each real server is set in configuration of real server pool. See • **Least Connection**: Forward a request to the real server with least connections (active connections and inactive connections) on. Actually, this algorithm also involves weight of the real servers when calculating. When weight of each real server in the pool is 1, calculation gives the result corresponding to the leastest connections. It varies when the weight changes. • **Fastest Response**: Forward a request to the real server with least active connections on. Like Least Connection, this involes weight of each real server as well.
IP Type	• **Use Interface IP Address** - Select to make the virtual server

Setting	Guidelines
	listen on the first IP address of the specified interface. • **Specify** - Select to expand the Address field to enter an IP address for the virtual server. Best practice is to specify another IP address if the virtual server is providing FTP service and current interface IP address is used for other services such as SNAT and FortiWAN GLB management.
Load Balance Persistence	Select a predefined or user-defined persistence configuration object. See Persistence - virtual server on page 128.
Connection Limit	Limit the number of concurrent connections. The default is 0 (disabled). The valid range is 1 to 1,048,576 concurrent connections. You can apply a connection limit per real server and per virtual server. Both limits are enforced. Attempted connections that are dropped by security rules are not counted. Not supported for FTP or SIP profiles.
Connection Rate Limit	With Layer 4 profiles, and with the Layer 2 TCP profile, you can limit the number of new connections per second. The default is 0 (disabled). The valid range is 1 to 86,400 connections per second. You can apply a connection rate limit per real server and per virtual server. Both limits are enforced. Attempted connections that are dropped by security rules are not counted. Not supported for FTP profiles.

- Click **Clone** at the right side of column to clone the configuration with a new name.
- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.

Log

The local log is a data store hosted on the FortiWAN system.

Typically, you use the local log to capture information about system health and system administration activities. We recommend that you use local logging during evaluation and verification of your initial deployment, and then configure remote logging to send logs to a log management repository where they can be stored long term and analyzed using preferred analytic tools.

Local log disk settings are configurable. You can select a subset of system events, traffic, and security logs.

FortiWAN-30E has no hard-disk embedded for local log database. Although FortiWAN-30E can store logs in RAM and report logs through the UI. Logs are removed, however, after system reboots.

The impact of logging on performance

- If you have a Log Server, store FortiWAN logs on the Log Server to avoid resource usage associated with writing logs to the local hard disk.
- If you don't need a traffic log, disable it to reduce the use of system resources.
- Reduce repetitive log messages. Use the alert email settings to define the interval that emails are sent if the same condition persists following the initial occurrence.
- Avoid recording log messages using low severity thresholds, such as information or notification, to the local hard disk for an extended period of time. Excessive logging frequency saps system resources and can cause undue wear on the hard disk and may cause premature failure.

Log Browsing

View logs that have been captured. You can filter your view and download complete logs.

Event log

The Event Log page shows logs related to system-wide status and administrator activity.

Use the following category filters to review logs of interest:

- **Configuration** - Configuration changes.
- **Admin** - Administrator actions.
- **Health Check** - Health check status changes. See Health checks on page 92.
- **System** - System operations, warnings, and errors.
- **LLB** - Notifications, such as bandwidth thresholds reached. See Link load balance (LLB) - Resources on page 113.
- **VS** - Notifications, such as connection limit reached. See About server load balancing (SLB) on page 245.
- **DNS** - Notifications, such as the status of associated local SLB and virtual servers. See Global Load Balancing on page 144.
- **Bandwidth Management** - Notifications, such as traffic matches bandwidth management policies. See Bandwidth on page 169.
- **Connection Limit** - Notification, such as connection limit is reached. See Connection limit on page 172.

Prerequisites

- Read-Write permission for Log & Report settings.

Access

- From the Dashboard, go to **Log > Log Browsing**, then select the **Event Log (default)** tab.

Settings

- Within each category, use Filter Setting controls to filter the table based on the values of matching data.

  Click **Download** to download the logs. Filters are applied to the set that's collected for download.
- You can filter the information within each log.
  a. Select a log type.
  b. Click **Filter Setting**, then choose an option from the drop-down list.
  c. Click **OK** to apply the filter. The list is updated.
- Click the page icon in the last column to see log the log settings.

Category Filters	Data Filters
Configuration	•Date

Category Filters	Data Filters
	• Time • Log Level • User • Message • Action
System, Admin	• Date • Time • Log Level • User • Message • Action • Status
Health Check	• Date • Time • Log Level • Group • Member • Module • Policy • Message
LLB (link load balancer)	• Date • Time • Log Level • Policy • Group • Member • Message • Status
DNS (Domain Name Server)	• Date

Category Filters	Data Filters
	• Time • Log Level • Policy • Message • Status
VS (Virtual Server)	• Date • Time • Log Level • Virtual Server • Group • Member • Message • Status
Bandwidth Management	• Date • Time • Log Level
Connection Limit	• Date • Time • Log Level

Sample Output

Event log - Configuration

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=15:50:37	Log time.
log_id	log_id=0000000085	Log ID.
log_level	log_level=information	Log Level

Column	Example	Description
msg_id	msg_id=522000	Message ID.
user	user=admin	User that performed the operation.
ui	ui=GUI(172.30.144.8)	User interface from which the operation was performed.
action	action=add	Administrator action: add, edit, delete.
cfgpath	cfgpath=firewall qos-queue	Configuration that was changed.
cfgobj	cfgobj=name	Configuration setting changed.
cfgattr	cfgattr=queue	Configuration value changed.
logdesc	logdesc=Change the configuration	Log description.
msg	msg=added a new entry 'queue' for "firewall qos-queue" on domain "root"	Log message.
type	type=event	Log type.
subtype	subtype=config	Log subtype.
vd	vd=root	Virtual domain.

Event log - System

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=15:50:37	Log time.
log_id	log_id=0003002078	Log ID.
log_level	log_level=alert	Log Level
msg_id	msg_id=44393	Message ID.
ui	ui=system	User interface from which the operation was performed.
status	status=success	Status of the performed operation.
logdesc	logdesc=Interface get a dynamic ip	Log description.
	msg=Interface port2 get a	
msg	dynamic ip 10.20.109.101/24	Log message.
	gateway 10.20.109.254	
	expire 1532248718	
type	type=event	Log type.
subtype	subtype=system	Log subtype.
vd	vd=root	Virtual domain.

Event log - Admin

Column	Example	Description
date	date=2014-12-01	Log date.
time	time=15:50:37	Log time.
log_id	log_id=0001002107	Log ID.

Column	Example	Description
log_level	log_level=information	Log Level
msg_id	msg_id=44344	Message ID.
user	user=admin	User that performed the operation.
ui	ui=GUI(10.12.102.66)	User interface from which the operation was performed.
action	action=login	Administrator action: login, logout.
status	status=success	Status of the action.
reason	reason=none	Reason for failed action.
logdesc	logdesc=Admin login	Log description.
msg	msg=User admin login successfully from GUI (10.12.102.66)	Log message.
type	type=event	Log type.
subtype	subtype=admin	Log subtype.
vd	vd=root	Virtual domain.

Event log - Health Check

Column	Example	Description

date	date=2014-12-01	Log date.
time	time=15:50:37	Log time.
log_id	log_id= 0002000730	Log ID.
log_level	log_level=alert	Log Level
msg_id	msg_id= 43577	Message ID.
modules	module=slb	Modules: LLB, slb, dns
policy	policy=none	LLB policy.
group	group=rsp-43-47	Real server pool (for module=slb)
member	member=2	Number of members in the group (for module=slb)
	msg=Pool name rsp-43-47	
	realserver name rsp-43-47, ip	
msg	192.168.63.47 and port 0	Log message.
	was detected as UP by	
	Health Check	
	LB_HLTHCK_ICMP	
action	action=health_check	Health check action.
type	type=event	Log type.
subtype	subtype=health_check	Log subtype.
vd	vd=root	Virtual domain.

Event log - LLB

Column	Example	Description
date	date=2018-07-22	Log date.
time	time=22:47:24	Log time.

Column	Example	Description
log_id	log_id=0004002277	Log ID.
log_level	log_level=notice	Log Level
msg_id	msg_id=435594	Message ID.
policy	policy=none	Undefined for LLB log.
group	group=none	Undefined for LLB log.
member	member=none	Undefined for LLB log.
status	status=success	Status of LLB action.
logdesc	logdesc=Link quality notification	Log description.
msg	msg=Link vpn1 rtt 707, jitter 45, packetloss 0	Log message.
action	action=prox_route	LLB action.
type	type=event	Log type.
subtype	subtype=llb	Log subtype.
vd	vd=root	Virtual domain.

Event log - DNS

Column	Example	Description
date	date=2018-07-16	Log date.
time	time=20:13:02	Log time.
log_id	log_id=0006002293	Log ID.
log_ level	log_level=alert	Log Level

msg_id	msg_id=435416	Message ID.
policy	policy=Unknown	Undefined for DNS logs.
status	status=failure	Success when remote server is online. Failure when remote server is offline.
logdesc	logdesc=GLB peer change state	Log description.
msg	msg=GLB Peer 10.20.127.101 is Disconnected	Log message.
action	action=none	Undefined for DNS logs.
type	type=event	Log type.
subtype	subtype=dns	Log subtype.
vd	vd=root	Virtual domain.

Event log - VS

Column	Example	Description
date	date=2018-07-16	Log date.
time	time=03:49:27	Log time.
log_id	log_id=0005000000	Log ID.
log_ level	log_level=warning	Log Level
msg_id	msg_id=43552	Message ID.
virtual_ server	virtual_server=vs2	The virtual server.
group	group=none	The real server pool.
member	member=none	Members of the real server pool.
status	status=failure	Success when connection or connection rate

		recovers from the limitation. Failure when connection or connection rate reaches the limitation.
logdesc	logdesc=Connection limit	Log description.
msg	msg=Virtual server vs2 is reached connection limit 10	Log message.
action	action=connection_limit	Virtual Server action.
type	type=event	Log type.
subtype	subtype=vs	Log subtype.
vd	vd=root	Virtual domain.

Event log - Bandwidth Management

Column	Example	Description
date	date= 2018-07-06	Log date.
time	time= 10:19:20	Log time.
log_id	log_id= 0007000000	Log ID.
log_level	log_level= information	Log Level
msg_id	msg_id= 44027	Message ID.
bm_policy_id	bm_policy_id=2	ID of the matched bandwidth management policy.
count	count=1	Count that the policy is matched.
type	type=event	Log type.
subtype	subtype=bm	Log subtype.
vd	vd=root	Virtual domain.

Event log - Connection Limit

Column	Example	Description

date	date= 2018-06-30	Log date.
time	time= 00:36:46	Log time.
log_id	log_id= 0008000000	Log ID.
log_level	log_level= information	Log Level
msg_id	msg_id= 43920	Message ID.
cl_policy_id	cl_policy_id=1	ID of the matched connection limit policy.
count	count=2	Count that the policy is matched.
type	type=event	Log type.
subtype	subtype=cl	Log subtype.
vd	vd=root	Virtual domain.

The value "none" appears in logs when the value is irrelevant to the status or action. For example, a health check log for a virtual server shows "none" in the Group and Member columns even though its real server pool and members are known because these details aren't relevant. Likewise, a health check log for a real server pool member shows "none" in the Policy column although its virtual server is known.

Security log

The security log automatically provides an aggregated view of security logs. There are two types of security logs:

- **Firewall** - Log for traffic matching firewall deny policies.
- **DoS** - Log for traffic encountering DoS.

Prerequisites

- Read-Write permission for Log & Report settings.

Access

- From the Dashboard, go to **Log > Log Browsing**, then select the **Security Log** tab.

Settings

- Within each category, use Filter Setting controls to filter the table based on the values of matching data.

  Click **Download** to download the logs. Filters are applied to the set that's collected for download.
- You can filter the information within each log.
  a. Select a log type.
  b. Click **Filter Setting**, then choose an option from the drop-down list.
  c. Click **OK** to apply the filter. The list is updated.
- Click the page icon in the last column to see log the log settings.

Sample Log Output

Security log - Firewall

Column	Example	Description
date	date= 2018-06-30	Log date.
time	time= 00:59:03	Log time.
log_id	log_id= 0201000000	Log ID.
log_level	log_level= notice	Log Level
msg_id	msg_id= 44057	Message ID.
firewall_policy_id	firewall_policy_id =1	ID of the matched firewall policy.
count	count=3	Count that the policy is matched.

Column	Example	Description
source	source= 192.168.63.47	Source IP of the denied packets.
source_country	source_country=reserved	Result mapping the source IP to GEO-IP database.
type	type= security	Log type.
subtype	subtype=firewall	Log subtype.
vd	vd=root	Virtual domain.

Security log - DoS

Column	Example	Description
date	date= 2018-07-22	Log date.
time	time= 01:48:13	Log time.
log_id	log_id= 0200000017	Log ID.
log_level	log_level= alert	Log Level
msg_id	msg_id= 492596	Message ID.
count	count=50	Count that packets being denied by the DOS prevent feature.

Column	Example	Description
severity	severity=high	
protocol	protocol=6	Protocol of the denied packets.
service	service=tcp	Service of the denied packets.
source	source=0.0.0.0	Source IP of he denied packets.
source_port	source_port=0	Source port of he denied packets.
destination	destination=10.20.108.101	Destination IP of he denied packets.
destination_port	destination_port=0	Destination port of he denied packets.
policy	policy=vs1	The virtual server suffering DoS attack.
action	action=deny	Action to the DoS packets.
source_country	source_country=reserved	Result mapping the source IP to GEO-IP database.

Column	Example	Description
destination_country	destination_country=reserved	Result mapping the destination IP to GEO-IP database.
type	type= security	Log type.
subtype	subtype=synflood	Log subtype.
vd	vd=root	Virtual domain.

Traffic log

The Traffic Log table displays logs related to traffic (sent and received bytes) served by the FortiWAN deployment.

Filter the following categories to see the logs:

- **Flow** - Logs for traffic transferred via LLB.
- **DNS** - Logs for traffic that DNS queries to GLB.
- **VS** - Logs for traffic of virtual servers.

Prerequisites

- Read-Write permission for Log & Report settings.

Access

- From the Dashboard, go to **Log > Log Browsing**, then select the **Traffic Log** tab.

Settings

- Within each category, use Filter Setting controls to filter the table based on the values of matching data.

  Click **Download** to download the logs. Filters are applied to the set that's collected for download.
- You can filter the information within each log.
  a. Select a log type.
  b. Click **Filter Setting**, then choose an option from the drop-down list.

c. Click **OK** to apply the filter. The list is updated.

- Click the page icon in the last column to see log the log settings.

The following information is available for all traffic log filters:

- Date
- Time
- Source
- Destination
- Service
- Trans Destination
- VS
- Real Server Name
- FQDN
- Resource IP
- Policy

Sample Log Output

Traffic log - Flow

Column	Example	Description
date	date= 2018-07-22	Log date.
time	time= 23:55:25	Log time.
log_id	log_id= 0100000000	Log ID.
log_level	log_level= information	Log Level
msg_id	msg_id= 435652	Message ID.
duration	duration=0	Flow duration.
received_bytes	received_bytes=73	Received bytes of the flow.
sent_bytes	sent_bytes=89	Sent bytes of the flow.
protocol	protocol=17	Protocol of the flow.

Column	Example	Description
source	source=192.168.63.45	Source IP of the flow.
source_port	source_port=49919	Source port of the flow.
destination	destination=8.8.8.8	Destination IP of the flow.
destination_port	destination_port=53	Destination port of the flow.
trans_source	trans_source=8.8.8.8	Trans source IP of the flow.
trans_source_port	trans_source_port=53	Trans source IP of the flow.
trans_destination	trans_destination=10.20.109.101	Trans destination IP of the flow.
trans_destination_port	trans_destination_port=49919	Trans destination port of the flow.
firewall_policy_id	firewall_policy_id=0	ID of firewall policy that the flow matches.
nat_policy_id	nat_policy_id=2	ID of NAT policy that the flow matches.
vip_policy_id	vip_policy_id=0	ID of VIP policy that the flow matches.
bm_policy_id	bm_policy_id=0	ID of Bandwidth management policy that the flow matches.
cl_policy_id	cl_policy_id=0	ID of connection limit policy that the flow matches.
llb_policy_id	llb_policy_id=0	ID of LLB policy that the flow matches.
source_country	source_country=reserved	Result mapping the

Column	Example	Description
		source IP to GEO-IP database.
destination_country	destination_country=United States	Result mapping the destination IP to GEO-IP database.
type	type=traffic	Log type.
subtype	subtype=flow	Log subtype.
vd	vd=root	Virtual domain.

Traffic log - DNS

Column	Example	Description
date	date= 2018-07-02	Log date.
time	time= 09:42:24	Log time.
log_id	log_id= 0101001523	Log ID.
log_level	log_level= information	Log Level
msg_id	msg_id= 44398	Message ID.
protocol	protocol=17	Protocol of the flow.
source	source= 192.168.60.50	Source IP of the flow.
source_port	source_port= 42130	Source port of the flow.
destination	destination= 192.168.63.1	Destination IP of the flow.
destination_port	destination_port=53	Destination port of the flow.
policy	policy=internal	Internal/external DNS query.
action	action=none	GLB action.

Column	Example	Description
fqdn	fqdn=www.aaa.com	FQDN for A/AAAA record query.
resource_ip	resource_ip=2.2.2.2	Answerred IP of the record.
source_country	source_country=reserved	Result mapping the source IP to GEO-IP database.
destination_ country	destination_country=reserved	Result mapping the destination IP to GEO-IP database.
type	type=traffic	Log type.
subtype	subtype=dns	Log subtype.
vd	vd=root	Virtual domain.

Traffic log - VS

Column	Example	Description
date	date= 2018-07-22	Log date.
time	time= 22:38:35	Log time.
log_id	log_id= 0102000000	Log ID.
log_level	log_level= information	Log Level
msg_id	msg_id= 435583	Message ID.
duration	duration= 100	Flow duration.
received_bytes	received_bytes= 280	Received bytes of the flow.
sent_bytes	sent_bytes= 434	Sent bytes of the flow.

Column	Example	Description
protocol	protocol= l6	Protocol of the flow.
source	Serviceftp	Sent bytes of the flow.
source	source=192.168.60.46	Protocol of the flow.
source_port	source_port=19707	Source IP of the flow.
destination	destination=10.20.108.101	Source port of the flow.
destination_port	destination_port=21	Destination IP of the flow.
trans_source	trans_source=192.168.60.46	Destination port of the flow.
trans_source_port	trans_source_port=19707	Trans source IP of the flow.
trans_destination	trans_destination=192.168.63.45	Trans source IP of the flow.
trans_destination_port	trans_destination_port=21	Trans destination IP of the flow.
real_server_name	real_server_name=pc45	Real server that virtual server forwarded the flow to.

Column	Example	Description
virtual_server	virtual_server=vs1-ftp	Virtual server handled the flow
action	action=none	Virtual server action.
source_country	source_country=reserved	Result mapping the source IP to GEO-IP database.
destination_country	destination_country=United States	Result mapping the destination IP to GEO-IP database.
type	type=traffic	Log type.
subtype	subtype=flow	Log subtype.
vd	vd=root	Virtual domain.

Log Setting

Log messages often help in determining the cause of a problem.

Depending on the type, log messages can appear in either the event, attack, or traffic logs. The FortiWAN appliance must be enabled to record event, attack, and traffic log messages; otherwise, you can't analyze the log messages for events of that type.

During troubleshooting, it can be useful to lower the logging severity threshold for more verbose logs, to include more information on less severe events.

Local log

Configure local logs to save specific information.

Access

- From the Dashboard, go to **Log > Log Settings**, then select the **Local Log** tab.

Prerequisites

- Read-Write permission for Log & Report settings.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

Setting	Guidelines
Status	Toggle **ON** to enable local logging.
File Size	Enter the maximum disk space used for a local log file. The default is 200 MB. When the current log file reaches this size, a new file is created.
Log Level	Choose the lowest severity to log: • **Emergency**—The system has become unstable. • **Alert**—Immediate action is required. • **Critical**—Functionality is affected. • **Error**—An error condition exists and functionality could be affected. • **Warning**—Functionality might be affected. • **Notification**—Information about normal events. • **Information**—General information about system operations. • **Debug**—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select **Error**, the system collects logs with levels Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with levels Alert and Emergency.
Disk Full	Select log behavior when the maximum disk space for local logs (30% of total disk space) is reached: • **Overwrite** - Continue logging. Overwrite the earliest logs.

Setting	Guidelines
	• **No Log** - Stop logging.
Event	Toggle **ON** to enable logging for events.
Event Category	This option is available when **Event** is on. Select the types of events to collect in the local log: • **Configuration**—Configuration changes. • **Admin**—Administrator actions. • **Health Check**—Health check status changes. • **System**—System operations, warnings, and errors. • **LLB**—Notifications, such as bandwidth thresholds reached. • **VS**—Notifications, such as connection limit reached. • **DNS**—Notifications, such as the status of associated local SLB and virtual servers. • **Bandwidth Management**—Traffic matches bandwidth management rule. • **Connection Limit**—Traffic matches connection limit rule.
Traffic	Toggle **ON** to enable logging for traffic processed by the load balancing modules.
Traffic Category	This option is available when **Traffic** is on. Select one. • **Flow** - Link Load Balancing traffic logs related to sessions and throughput. • **VS** - Virtual server traffic logs related to sessions and throughput. • **DNS** - Global DNS server traffic logs related to DNS requests.
Security	Toggle **ON** to enable logging for traffic processed by the security modules.
Security Category	• **DoS** - SYN flood protection logs. • **Firewall** - Packets are denied by firewall rules.

Click **Refresh** to reset any modifications you made since the last **Save**.

Remote log

A remote syslog server collects logs for long-term storage. The logs can be used with your preferred analytic tools.

Access

- From the Dashboard, go to **Log > Log Setting**, then select the **Remote Log** tab.

Prerequisites

- Read-Write permission for Log & Report settings.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Status	Toggle **ON** to enable remote logging.
Address	Enter the IP address of the syslog server.
Port	Enter the listening port number of the syslog server. Usually this is UDP port 514.
Log Level	Choose the lowest severity to log: • **Emergency**—The system has become unstable. • **Alert**—Immediate action is required. • **Critical**—Functionality is affected. • **Error**—An error condition exists and functionality could be affected. • **Warning**—Functionality might be affected. • **Notification**—Information about normal events. • **Information**—General information about system operations. • **Debug**—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select **Error**, the system collects logs with levels Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with levels Alert and Emergency.
CSV	Toggle **ON** to save logs in CSV format.
Facility	Choose an identifier that isn't used by any other device on your network when sending logs to syslog.
Event	Toggle **ON** to enable logging for events.
Event Category	This option is available when **Event** is on. Select the types of events to collect in the local log: • **Configuration**—Configuration changes. • **Admin**—Administrator actions. • **Health Check**—Health check status changes. • **System**—System operations, warnings, and errors.

Setting	Guidelines
	• **LLB**—Notifications, such as bandwidth thresholds reached. • **VS**—Notifications, such as connection limit reached. • **DNS**—Notifications, such as the status of associated local SLB and virtual servers. • **Bandwidth Management**—Traffic matches bandwidth management rule. • **Connection Limit**—Traffic matches connection limit rule.
Traffic	Toggle **ON** to enable logging for traffic processed by the load balancing modules.
Traffic Category	Available when Traffic is on, choose at least one option. • **Flow** - Link Load Balancing traffic logs related to sessions and throughput. • **VS** - Virtual server traffic logs related to sessions and throughput. • **DNS** - Global DNS server traffic logs related to DNS requests.
Security	Select to enable logging for traffic processed by the security modules.
Security Category	Available when Security is on, choose at least one option. • **Firewall** - Packets are denied by filewall rules. • **DoS** - SYN flood protection logs.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Alert Email

Set up email alerts to let you know when critical events happen. Also, use alerts instead of logging to improve performance.

Prerequisites

- Read-Write permission for Log & Report settings.

Access

- From the Dashboard, go to **Log > Log Setting**, then select the **Alert Email** tab.

Settings

- Set your options as needed, then **Save**.

Setting	Guidelines
By category	Toggle ON to enable receiving emails by category.
Categories	This option is available when **By category** is on. Choose the ones you want. • ha • admin • config • diskfull • healthcheck • cert
Loglevel	This option is available when **By category** is off. Choose the lowest severity to log: • **Emergency**—The system has become unstable. • **Alert**—Immediate action is required. • **Critical**—Functionality is affected. • **Error**—An error condition exists and functionality could be affected. • **Warning**—Functionality might be affected. • **Notification**—Information about normal events. • **Information**—General information about system operations. • **Debug**—Detailed information about the system that can be used to troubleshoot unexpected behavior. For example, if you select **Error**, the system collects logs with levels Error, Critical, Alert, and Emergency. If you select **Alert**, the system collects logs with levels Alert and Emergency.
Interval	Enter the number of minutes to receive an email.
From	Enter the email address you want the emails to appear to be from. This is useful if you want to assign a rule to incoming emails based on From address.

Recipient

Create recipients to receive the email alerts.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.
- Click **Add** to open the configuration editor.

Setting	Guidelines
Name	Enter the recipient name.
Mail to	Enter the recipient email address.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.

Diagnostic

These topics describe FortiWAN diagnostic tools.

Packet capture

Prerequisites

- A good understanding of tcpdump and filter expressions. See http://www.tcpdump.org/manpages/pcap-filter.7.html.
- Read-Write permission for System settings.

Access

- From the Dashboard, go to **Diagnostic > Packet Capture**.

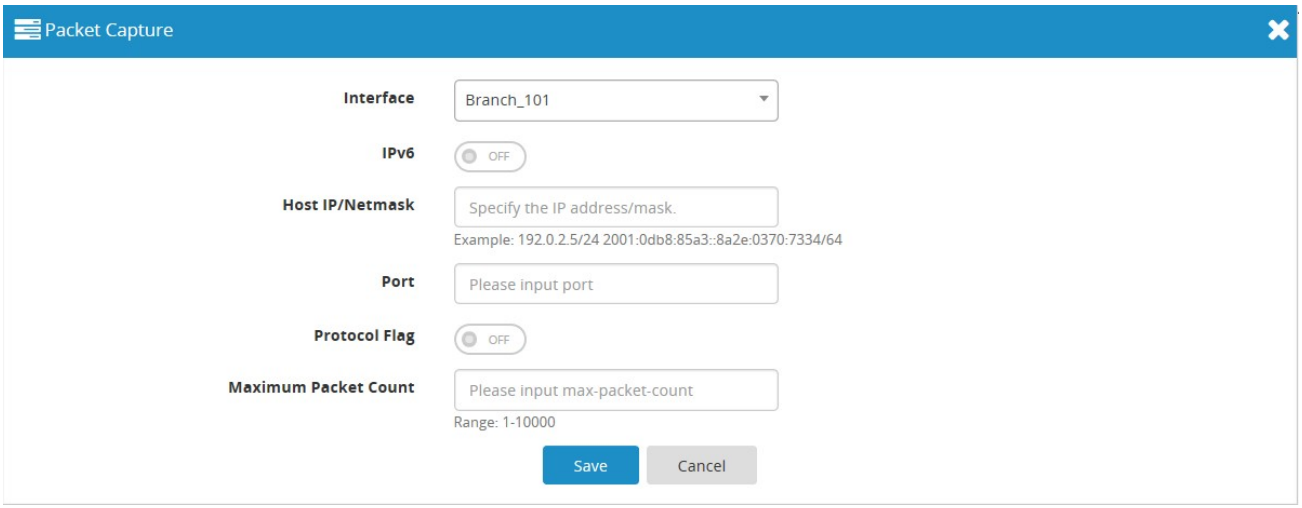

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

- Click **Add** to open the configuration editor.

Setting	Guidelines
Interface	Choose an interface from the drop down list. See Interface settings on page 63.
IPv6	Toggle **ON** to enable IPv6.
Host IP/Netmask	Enter the IP address/mask. For example, 192.0.2.5/24 2001:0db8:85a3::8a2e:0370:7334/64.
Port	Enter the port number.
Protocol Flag	Toggle **ON** to enable the Protocol Flag, then select the protocols you want. •arp •tcp •udp •icmp
Maximum Packet Count	Enter the number of maximum packets you want to allow. The range is 1 to 10000.

- To **modify**, double-click the item from the list, or click the **Pencil** icon at the end of the row.
- To **delete**, check the box for the item you want to remove, then click **Delete** at the top of the page, or click the **Delete** icon at the end of the row.
- Click **Clone** at the right side of column to clone the configuration with a new name.
- Click **Run** ▶ to start the packet capture.
- Click **Stop** ■ to stop the packet capture.
- Click **Download** to save the packet capture to your local drive in .CAP format.

`tcpdump` can also be used from the Console. See

Reducing the impact of packet capture on system performance

Packet capture can be useful for troubleshooting but can be resource intensive. To minimize the impact on system performance, use packet capture only during periods of minimal traffic. Use a local

console CLI connection rather than a Telnet or SSH CLI connection, and be sure to stop the command when you are finished.

Diff

Compare backups of the core configuration file with your current configuration. This can be useful if:

- A previously configured feature is no longer functioning, and you are not sure what in the configuration has changed.
- You want to recreate something configured previously, but don't remember what the settings were.

Difference-finding programs, such as WinMerge, can help you quickly compare various versions of your output and easily find differences. They can compare your configurations, line by line, and highlight parts that are new, modified, or deleted.

Figure 20 - Configuration differences highlighted in WinMerge

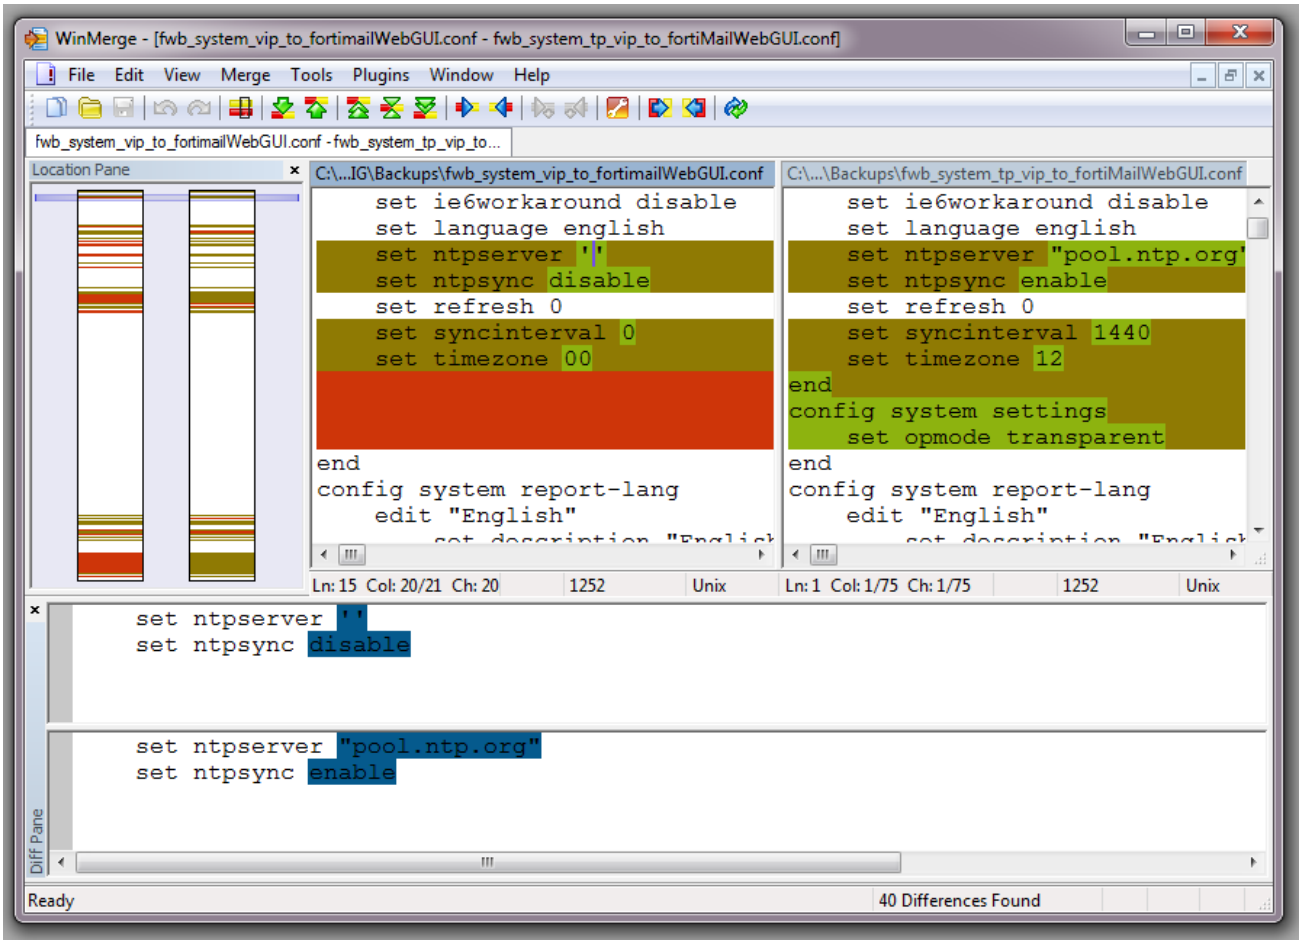

For instructions, see the documentation for your diff program.

Console

You can use the FortiWAN Console to enter CLI commands.

- From the top of the UI, click **Console** Console.
  The console window opens. Type your commands.
- To close the console window, click **Console** Console again.

You can also access the console by telnet or SSH.

Access control CLI commands

For each config command, there is an equivalent get/show command.

- The config commands require write permission.
- The get/show commands require read permission.

UI Menus	CLI Commands
System	config system diagnose hardware diagnose sniffer diagnose system execute date execute ping execute ping-options execute
Networking	config router config bandwidth-management config connection-limit config nat config vpn
Server Load Balance	config server-load-balance
Link Load Balance	config link-load-balance
Global DNS Server	config global-dns-server execute discovery-glb-virtual-server

UI Menus	CLI Commands
Share Resource	config share-resource
Log	config log execute log execute rebuild

You can also use CLI commands directly to the device using an SSH client such as PuTTY.

CLI: execute commands

You can use the command-line interface (CLI) execute commands to run system management utilities, such as backups, upgrades and reboots; and network diagnostic utilities, such as `nslookup`, `ping`, `traceroute`, and `tcpdump`.

Following is a list of execute commands:

`FortiWAN-VM # execute ?`	
`backup`	backup
`caching`	caching management
`certificate`	certificate
`checklogdisk`	find and auto correct errors on the log disk
`clean`	clean
`config-sync`	config sync
`date`	set/get date and time
`discovery-glb-vir-tual-server`	Sync virtual servers from glb server, add them to the virtual server list

dns-caching	dns caching management
dnssec	dnssec files management
dumpsystem	dump system information for debugging purpose
dumpsystem-file	dumpsystem-file
expanddatadisk	repartition data disk
factoryreset	reset to factory default
fixlogdisk	correct errors on the log disk
formatlogdisk	format log disk to enhance performance
geolookup	lookup geography information for IP address
glb-clear	Clear GLB information
glb-dprox-lookup	lookup GLB dynamic proximity information
glb-persistence-lookup	lookup GLB persistence information
ha	ha
health-check-verify	health check verify
isplookup	lookup ISP name and isp-address for IP address
log	log management
nslookup	nslookup
packet-capture	packet-capture <Port Number> [filter] (Only IPv4)
packet-capture-file	packet-capture-file
packet-capture6	packet-capture6 <Port Number> [filter] (Include IPv6)
ping	ping <host name | host ip>
ping-option	ping option settings
ping6	ping <host name | host ipv6>

ping6-option	ping6 option settings
reboot	reboot the system
reload	reload appliance
restore	restore
saml-idp	saml-idp
shutdown	shutdown appliance
ssh	Simple SSH client
ssl-forward-proxy-certificate-caching	ssl-forward-proxy-certificate-caching management
statistics-db	statistics db management
telnet	Simple telnet client
traceroute	traceroute
vm	vm
web-category-test	Test a url find its web-category

CLI: diagnose commands

Use the CLI diagnose commands to gather diagnostic information that is useful to Fortinet Customer Care when diagnosing any issues with your system. The commands are similar to the Linux commands used for debugging hardware, system, and IP networking issues.

The most important command for customers to know is `diagnose debug report`. This prepares a report you can give to your Fortinet support contact to assist in debugging an issue.

The following shows the diagnose commands:

FortiWAN-VM # diagnose ?	
debug	debug
hardware	hardware
llb	llb

netlink	netlink
server-load-balance	server-load-balance
share-resource	share-resource
sniffer	sniffer
soft-switch	soft-switch
system	system

CLI: System dump

The system includes utilities for generating system dump files that can help Fortinet support engineers analyze an issue for you. At present one can use CLI to generate and upload dump files:

Dump kernel and user space information when the system is still responsive

```
FortiWAN-VM # execute dumpsystem
This operation will reboot the system! Do you want to continue? (y/n)y
Begins to dump userspace information Begins to dump kernel information

FortiWAN-VM # execute dumpsystem-file list
-rw------- 1 0 0 96719189 Mar 15 13:35coredump-2016-03-15-13_35
-rw-r--r-- 1 0 0 16654391 Mar 15 13:34user_coredump_2016_03_15_13_34_
46.tar.bz2

FortiWAN-VM # execute dumpsystem-file upload tftp coredump-2016-03-15-
13_35 172.30.184.77 coredump-2016- 03-15- 7% |** | 7152k 0:09:58ETA
```

CLI: Change Admin password

```
FortiWAN-VM # config system
admin FortiWAN-VM (admin) #
edit admin
FortiWAN-VM (admin) # set password
<string> Current password for 'admin':
FortiWAN-VM (admin) # end
```

CLI: Connect to the UI

You can connect to the UI from the device.

1. Use an SSH client such as PuTTY to make an SSH connection to 192.168.1.99 (tcp port 22).
2. Acknowledge any warnings and verify and accept the FortiWAN SSH key.
3. Enter the user name **admin** and no password.
4. Use the following command sequence to configure the management interface:

```
config system interface
   edit
   <interface_name>
      set ip
      <ip&netmask>
   set allowaccess {http https ping snmp ssh telnet}
   end
end
```

The system processes the update and disconnects your SSH session because the UI has a new IP address. At this point, you should be able to connect to the CLI from a host on the management subnet you just configured. You can verify the configuration remotely.

You can also access the UI from a browser. See Step 2: Configure the management interface on page 8.

CLI: Configure a VLAN interface

```
config system interface edit
   <specified_name> set type vlan
   set vlanid <number>
   set interface <port_name>
      set ip
   <ip&netmask> end
```

You can also perform this action in the UI. See Example: How to configure a VLAN interface on page 68.

CLI: Configure a software switch interface

```
config system interface edit
   <specified_name>
   set type soft-switch set vlanid <number>
      set member <port_name> <port_name>
      set ip
   <ip&netmask> end
```

CLI: Set System Time

Configure system time using NTP

```
config system time
   ntp set ntpsync
   enable
   set ntpserver {<server_fqdn> |
   <server_ipv4>} set syncinterval
   <minutes_int>
end
```

Configure system time manually

```
config system time
   manualset zone
   <timezone_index>
   set daylight-saving-time {enable | disable} end
   execute date <MM/DD/YY> <HH:MM:SS>
```

CLI: Configure DNS

```
config system dns
set primary
<address_ipv4>set
secondary <address_ipv4>
end
```

CLI: Static route settings

```
config router static edit 1
   set destination <ip address/netmask> set gateway <ip address>
   set distance <value>
end
```

You can also perform this action in the UI. See Static route settings on page 80.

CLI Configure Network Interfaces

Show network interface addresses and subnets

```
show system interface
```

Enable an interface

```
config system interface
   edit port2
      set status up
end
```

Configure a physical interface

```
config system interface
   edit
   <interface_name>
      set ip <ip&netmask>
      set allowaccess {http https ping snmp ssh telnet}
   end
end
```

Configure an aggregate interface

```
config system interface
```

```
    edit
    <specified_name>
    set type agg
    set aggregate-mode {802.3ad | balance-alb | balance- rr | balance-tlb |
    balance-xor | broadcast}
    set member <port_name>
       <port_name>
    set ip
       <ip&netmask>
end
```

CLI Disable Health Check

```
config link-load-balance link edit <link_name>
   set health-check-ctrl <disable | enable> next
end
```

To perform this action in the UI, simply delete the Health Check item.
See Health checks on page 92.

CLI: Update firmware

The CLI upgrade procedure replaces the firmware on the alternate partition and reboots. The alternate (upgraded) partition becomes the active, and the active becomes the alternate.

The CLI does not have an equivalent of the UI Boot Alternative Firmware command.

Prerequisites

- Read the release notes for the version you plan to install. The release notes contain the latest information and supersede anything in this documentation.
- Be able to use TFTP to transfer the firmware file to the FortiWAN. Download and install a TFTP server, such as tftpd (Windows, Mac OSX, or Linux), on a server on the same subnet as the FortiWAN.

- Download the firmware file from the Fortinet Customer Service & Support website: https://support.fortinet.com/.
- Copy the firmware image file to the root directory of the TFTP server.
- Back up your configuration before beginning this procedure. See Back up and restore on page 42.
- Have superuser permission (user admin) to upgrade firmware.

TFTP isn't secure, and it doesn’t support authentication. You should run it only on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn off tftpd off immediately after completing this procedure.

To install firmware

1. Connect your management computer to the FortiWAN console port using an RJ-45-to-DB-9 serial cable or a null- modem cable.
2. Initiate a connection to the CLI and log in as the user **admin**.
3. Use an Ethernet cable to connect FortiWAN port1 to the TFTP server directly, or connect it to the same subnet as the TFTP server.
4. If necessary, start the TFTP server.
5. Use the following command to transfer the firmware image to the FortiWAN system:

   ```
   execute restore image tftp <filename> <tftp_ipv4>
   ```

   Example upgrade

   ```
   FortiWAN-VM # execute restore image tftp FWN_VM-v501-build0166-
   FORTINET.out
   192.0.2.1 This operation will replace the current firmware version!
   Do you want to continue? (y/n)y
   Connect to tftp server
   192.0.2.1 ... Please wait...
   ###################################################################
   Get image from tftp server OK. Check image trailer OK. Check image
   OK.
   FortiWAN-VM
   #
   ```

   Example downgrade

   ```
   FortiWAN-VM # execute restore image tftp FWN_VM-v501-build0165-
   FORTINET.out
   ```

```
192.0.2.1 This operation will replace the current firmware version!
Do you want to continue? (y/n)y
Connect to tftp server
192.0.2.1 ... Please wait...
############################################################
Get image from tftp server OK. Check image trailer OK.
This operation will downgrade the current firmware version! Do you
want to continue? (y/n)y
FortiWAN-VM #
```

6. To verify the upgrade, display the system version number:

```
FortiWAN-VM # get system status
Version: FortiWAN-VM v5.0.1,build0166,180722
VM Registration: Valid: License has been successfully authenticated
with registration servers.
VM License File: License file and resources are valid.
VM Resources: 1 CPU/1 allowed, 1620 MB RAM/2048 MB allowed, 23 GB
Disk/1024 GB allowed
...
```

If the download fails after the integrity check with the error message `invalid compressed format (err=1`, but the firmware matches the integrity checksum on the Fortinet Customer Service & Support website, try a different TFTP server.

You can also perform this action in the UI. See Updating firmware on page 40.

Reverse path route caching

By default, reverse path route caching is enabled. FortiWAN caches a reverse path route for inbound traffic so itcan forward reply packets to the ISP link that forwarded the corresponding request packet. This is useful when your site receives traffic from multiple ISP links. In the example below, the reverse path pointer ensures that client traffic received from ISP1 is returned through ISP1.

Figure 21 - Reverse path route caching enabled

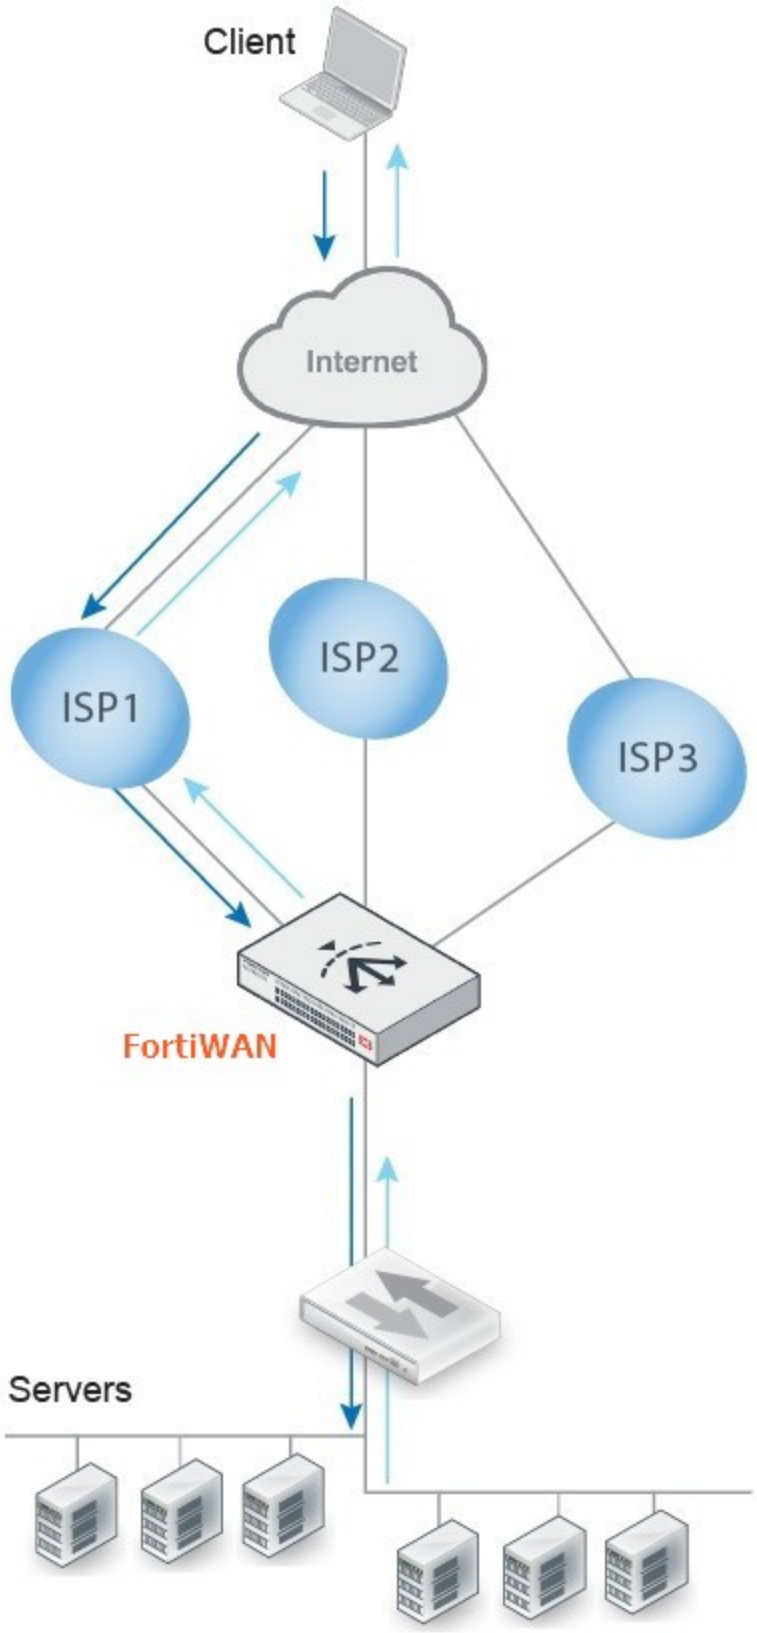
Client
Internet
ISP1
ISP2
ISP3
FortiWAN
Servers

When reverse path caching isn't enabled, the system forwards reply packets based on the results of routing lookup.

- To enable/disable reverse path route caching, use the `config router setting` CLI command:

```
FortiWAN-VM # config router
setting FortiWAN-VM (setting) #
get
rt-cache-strict : disable
rt-cache-reverse : enable
ip-forward : enable
ip6-forward : enable
icmp-redirect-send : disable
FortiWAN-VM (setting) # set rt-cache-reverse
disable FortiWAN-VM (setting) # end
FortiWAN-VM # get router
setting rt-cache-strict : disable
rt-cache-reverse : disable
ip-forward : enable
ip6-forward : enable
icmp-redirect-send : disable
```

The `rt-cache-strict` option is disabled by default. Enable it when you want to send reply packets only via the same interface that received the request packets. When enabled, source interface becomes part of the matching tuple that FortiWAN uses to identify sessions, so reply traffic is forwarded from the same interface thatreceived the traffic. (Normally each session is identified by a 5-tuple: source IP, destination IP, protocol, sourceport, and destination port.)

If the `rt-cache-reverse` option is enabled, you can use the `config rt-cache-reverse-exception` command to maintain an exceptions list for source IP addresses that should be handled differently. For example, if you configure an exception for 192.168.1.0/24, FortiWAN will not maintain a pointer to the ISP for traffic from source 192.168.1.18. Reply packets will be forwarded based on the results of routing look up.

```
FortiWAN-docs # config router
setting FortiWAN-docs (setting) #
get
rt-cache-strict : disable
rt-cache- reverse : enable
ip- forward : enable
ip6-forward : enable
icmp-redirect-send : disable
FortiWAN-docs (setting) # config rt-cache-reverse- exception
FortiWAN-docs (rt-cache-rever~e) # edit 1
```

```
Add new entry '1' for node 3740
FortiWAN-docs (1) # set ip-netmask 192.168.1.0/24
FortiWAN-docs (1) # end
FortiWAN-docs (setting) # end
```

About server load balancing (SLB)

FortiWAN's Server Load Balancing routes traffic to available destination servers based on health checks and load-balancing algorithms. SLB improves application availability and performance, which directly improves user experience.

The physical distance between clients and the servers in your backend server farm has a significant impact on server response times. Besides physical distance, the most important factors contributing to server performance are:

- Number of simultaneous connections and requests that the servers can handle
- Load distribution among the servers

The purpose of an SLB is to give you multiple methods for optimizing server response times and server capacity.

After you have deployed an SLB, traffic is routed to the SLB *virtual* server instead of the destination *real* servers.

The following figure shows an example of a basic load balancing deployment. The FortiWAN appliance is deployed in front of a server farm, and the network interfaces are connected to three subnets: a subnet for management traffic; a subnet that hosts real servers A, B, and C; and a different subnet that hosts real servers D, E, and F. The FortiWAN system performs health checks on the real servers and distributes traffic to them based on system logic and user-defined settings.

Figure 22 - Basic network topology

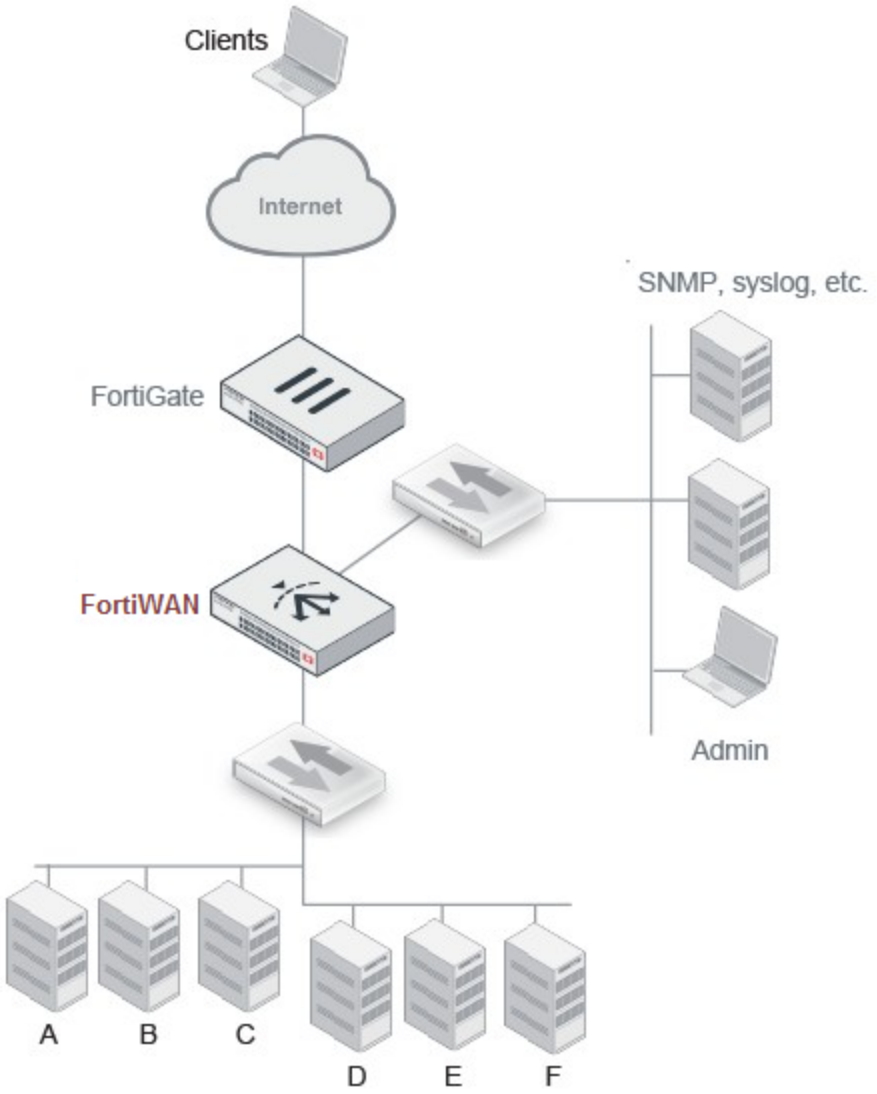

The configuration object framework supports the granularity of FortiWAN application delivery control rules. You can configure specific options and rules for one particular type of traffic, and different options and rules for another type.

Workflow

1. Configure real servers.
   Real servers are the backend servers providing network services.
2. Configure health check rules (optional).
   In many cases, you can use predefined health check rules.

3. Configure server pools.

   Server pools are the backend servers you want to load balance and specify the health checks used to determine server availability.

4. Configure persistence rules, profile components, and NAT pools.

   You can skip this step if you want to select from predefined persistence rules, profiles, and run the virtual server without using FullNAT for forwarding method.

5. Configure the virtual server.

   When you configure a virtual server, you select from predefined and custom configuration objects.

About global load balancing

In a global load balancing deployment, you configure DNS server and global load balancing details only on the global FortiWAN instance. The configuration framework enables granular administration and fine tuning of both the DNS server and the global load balancing framework.

The following shows the basic configuration elements for global load balancing and the recommended order for creating the configuration objects. The order is important for initial configurations because complex configuration elements like policies often include references to simple configuration objects like the remote DNS servers(forwarders) or DNS64 rules, but the simple elements must be created first.

Figure 23 - Global load balancing configuration summary

Workflow - DNS server

1. Complete the zone configuration. The global load balancing framework generates the zone configuration for zones that include the FortiWAN virtual servers.
2. Configure general DNS settings. Complete the zone configuration.
3. Configure remote DNS servers (forwarders) that you might reference in the zone configuration.

Workflow - Global load balancing

1. Create the data center, servers, virtual server pool, and host configurations that are the framework for associating locations with virtual servers and generating the DNS zone configuration

and resource records. You can adjust the dynamic proximity and persistence settings at any time.

2. Review the generated DNS zone configuration.

About High Availability

FortiWAN appliances can be deployed as standalone units or as high availability (HA) clusters.

A **cluster** is two or more nodes. A **node** is an instance of the appliance or system. In a cluster, one node is the primary node, also called the master node. The other members of the cluster are secondary nodes, also called child nodes.

The primary node has a special role. It has a one-to-many relationship with member nodes. Both configuration updates and software updates are initiated by the primary node and pushed to member nodes.

The system selects the primary node based on the following criteria:

- Link health (if monitor ports links are down, the node is considered down)
- Remote IP monitor health check results
- Override setting (prefers priority to up time)
- Most available ports
- Highest up time value
- Lowest device priority number (1 has greater priority than 2)
- Highest-sorting serial number - Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

HA solutions depend on two types of communication among cluster members:

- **Synchronization** - During initialization, the primary node pushes its configuration (with noted exceptions) to member nodes.
- **Heartbeats** - A cluster node indicates to other nodes in the cluster that it is up and available. The absence of heartbeat traffic indicates the node isn't up and is unavailable.

When a cluster is running, it runs in Active-Passive mode. Only the primary node is active, so it’s the only node that receives traffic from adjacent routers. Typically, there’s one other node that’s in standby mode. It assumes active status if the primary node undergoes maintenance or otherwise becomes unavailable.

HA cluster topology

The following shows an active-passive cluster in a single network path. In an active-passive cluster, the primary node is the active node that handles all traffic. In the event that the primary node experiences hardware failure or system maintenance, fail over takes place. In fail over, the standby

node becomes the primary node and processes the traffic that's forwarded along the network path. The new primary node sends gratuitous ARP to notify the network to direct traffic for the virtual MAC addresses (vMAC) to its network interfaces. It takes the IP addresses of the unresponsive node.

Figure 24 - Basic active-passive cluster

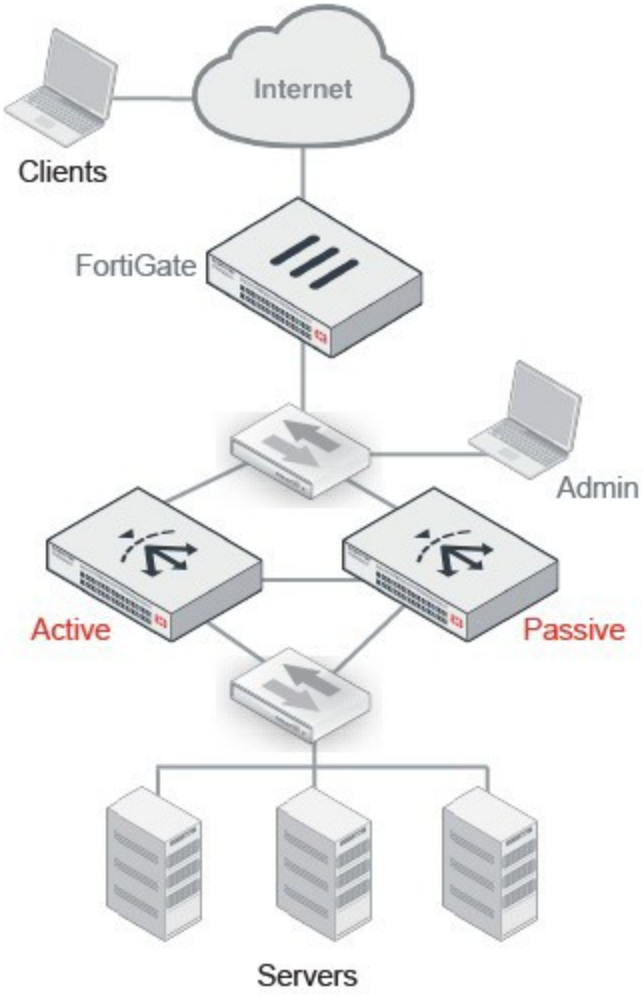

The following figure hows an active-passive cluster in a redundant path. A topology like this is a best practice because it is fully redundant, with no single point of failure. If the gateway, load balancer, or switch were to fail, the fail over path is chosen.

Figure 25 - Redundant path active-passive cluster

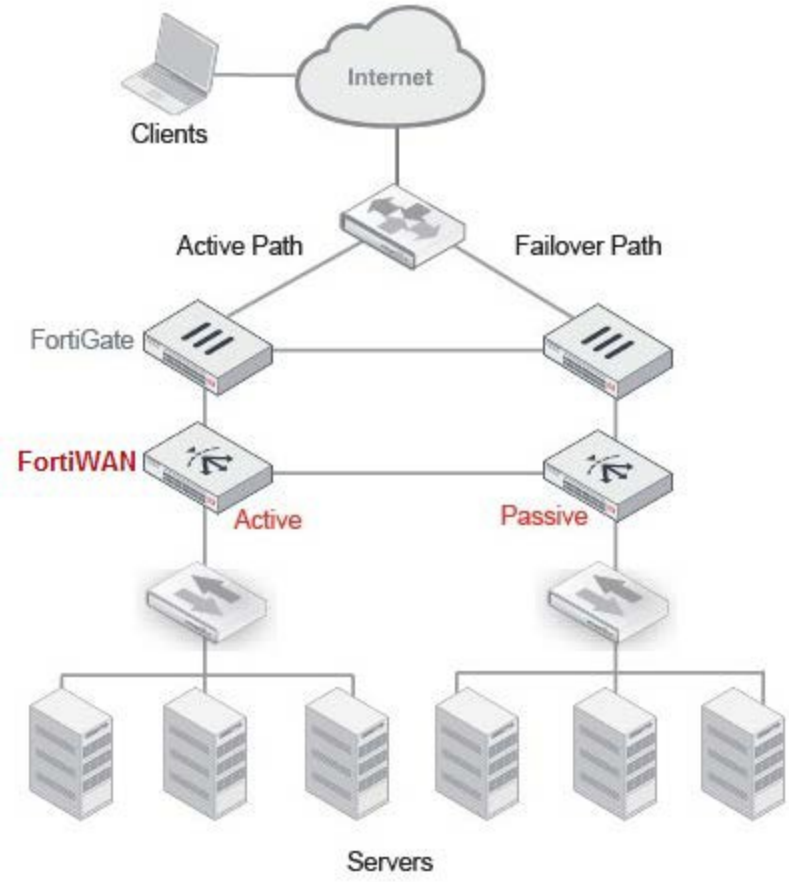

HA synchronization

The master node pushes most of its configuration to the other member nodes. This is known as synchronization. If automatic synchronization is enabled, synchronization occurs immediately when an appliance joins the cluster, and thereafter every 30 seconds. If synchronization isn't enabled, you must initiate synchronization manually.

Synchronization includes:

- Core CLI-style configuration file
- Health check status

For most settings, you configure only the primary node, and its settings are pushed to other members.

The following settings are not synchronized. All other settings are synchronized.

Setting	Guidelines
host name	The host names are not synchronized to enable you to use unique names.
SNMP system information	Each member node has its own SNMP system information so that you can maintain accurate, separate data in SNMP collections. However, the network interfaces of a standby node are not active, so they can't be actively monitored with SNMP.
RAID level	RAID settings are hardware-dependent and determined at boot time by looking at the drives (for software RAID) or the controller (hardware RAID), and are not stored in the system con- figuration. Therefore, they are not synchronized.
HA settings	Most of the HA configuration isn't synchronized in order to support HA system operations. In particular: • **Priority and Override** settings - These settings are used to select a primary node, so they are not synchronized to enable dif- ferentiation. • **Group ID** - Nodes with the same Group ID join a cluster. The set- ting precedes and determines group membership, so it is set manually. • **HA mode** - Many administrators prefer to be able to switch the primary node from an HA mode to standalone mode without the other nodes following suit, or to switch a secondary node to stan- dalone mode and have that setting not overwritten by periodic synchronization, so the HA mode setting isn't pushed from the primary node to the member nodes.

In addition to the HA configuration, some data is also not synchronized:

- Log messages - These describe events that happened on that specific appliance. After a fail over, you might notice that there is a gap in the original active appliance log files that cor- responds to the period of its down time. Log messages created during the time when the standby was acting as the active appliance (if you have configured local log storage) are stored there, on the original standby appliance.
- Generated reports (statistics in Dashboard) - Like the log messages that they are based upon, reports also describe events that happened on that specific appliance.

You can view the status of cluster members from the dashboard of the primary node. You might need to log into the system for the a non-primary member node (See) in the following situations:

- To configure settings that are not synchronized.
- To view log messages recorded about the member node itself on its own hard disk.
- To view traffic reports for traffic processed by the member node.

Monitoring an HA cluster

To monitor HA information, go to **Dashboard > Main**.

Figure 26 - HA Information widget in the dashboard

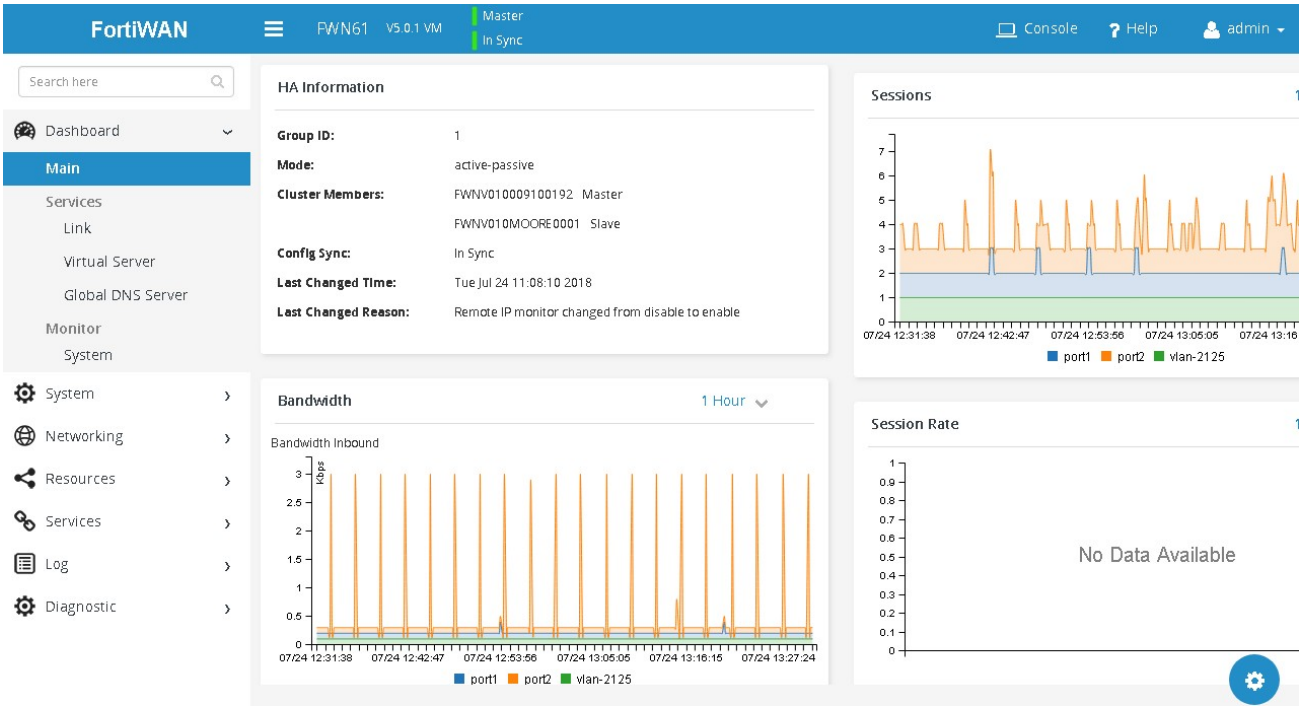

You can use also use log messages, alert emails, and SNMP to monitor HA events, such as when fail over has occurred. The system logs HA node status changes as follows:

- When HA is initialized: `HAdeviceInit`
- When a member joins a group: `Member(FWN2HE5118000002)jointotheHAgroup`
- When the HA configuration is changed from standalone to an active-passive cluster mode: `HAdeviceintoSlavemode`

The following figure shows FortiWAN HA event objects in log browsing.

Figure 27 - FortiWAN HA event objects in GUI log browsing

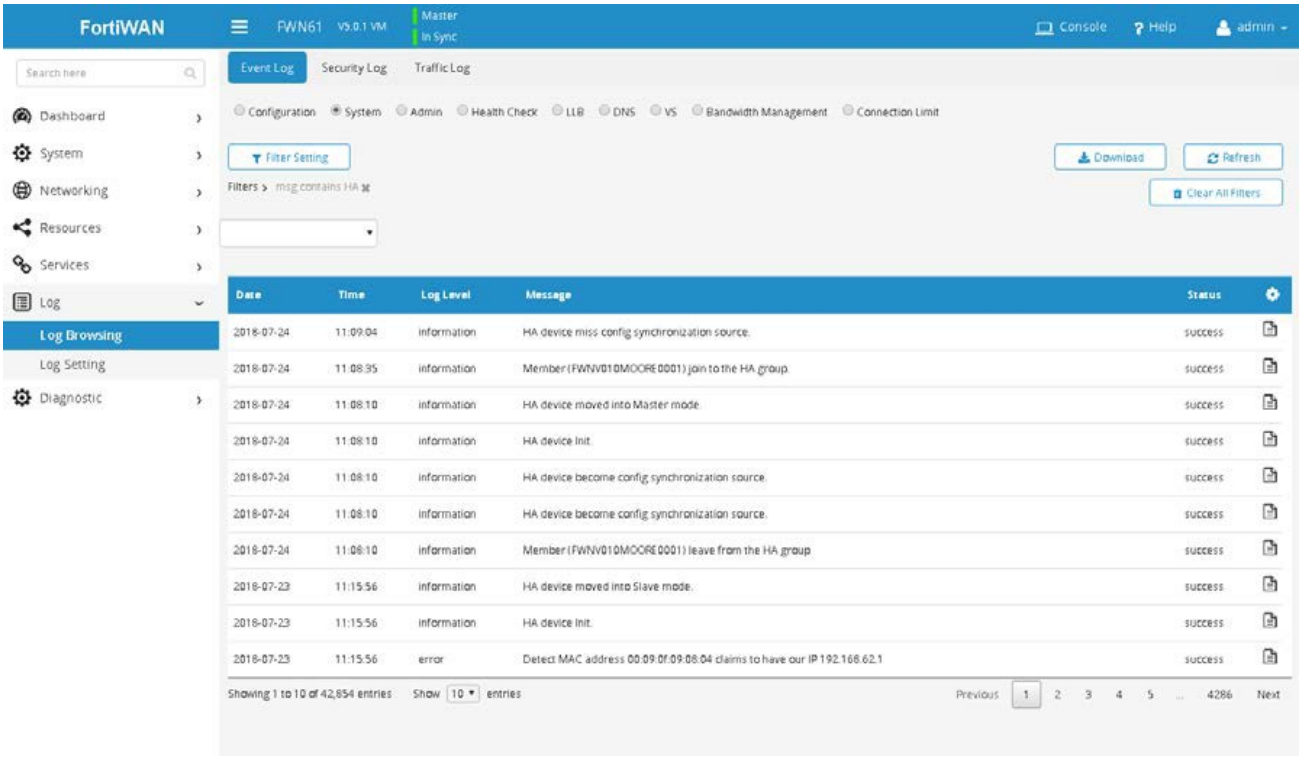

Updating firmware for an HA cluster

You can upgrade firmware on all nodes in a cluster from the primary node. The following process occurs when you perform the HA upgrade procedure:

1. The primary node pushes the firmware image to the member nodes.
2. The primary node notifies the member nodes of the upgrade, and it takes their user traffic during the upgrade.
3. The upgrade command is run on the member nodes, the systems are rebooted, and the member nodes send the primary node an acknowledgment that upgrade has been completed.
4. The upgrade command is run on the primary node, and it reboots. When the system is rebooting, a member node assumes primary status, and the traffic fails over from the former primary node to the new primary node.

After the upgrade process is completed, the system determines whether the original node becomes the primary node, according to the HA Override setting:

- If Override is enabled, the cluster considers the Device Priority setting. Both nodes usually make a second failover in order to resume their original roles.

- If Override is disabled, the cluster considers uptime first. The original primary node will have a smaller uptime due to the order of reboots during the firmware upgrade. Therefore it will not resume its active role; instead, the node with the greatest uptime will remain the new primary node. A second failover will not occur.

Reboot times vary by the appliance model, and also by differences between the original firmware version and the firmware version you are installing.

The administrator procedure for an HA cluster is similar to the procedure for installing firmware on a standalone appliance. To ensure minimal interruption of service to clients, use the following steps. The same procedure applies to both active-active and active-passive clusters.

If downgrading to a previous version, don't use this procedure. The HA daemon on a member node might detect that the primary node has older firmware, and attempt to upgrade it to bring it into sync, undoing your downgrade.

Instead, switch out of HA, downgrade each node individually, then switch them back into HA mode.

Prerequisites

- Download the firmware file from the Fortinet Customer Service & Support website: https://support.fortinet.com/.
- Read the release notes for the version you plan to install.
- Back up your configuration before beginning this procedure. Reverting to an earlier firmware version could reset settings that are not compatible with the new firmware.
- You must have super user permission (user admin) to upgrade firmware.
- Verify that the cluster node members are powered on and available on all of the network interfaces that you have configured. If required ports are not available, HA port monitoring could inadvertently trigger an additional failover, resulting in traffic interruption during the firmware update.

Access

- Log into the UI of the **primary** node as the `admin` administrator.
- From the Dashboard, go to **System > Settings**, then select the **Maintenance** tab.

Settings

- Set your options as needed, then **Save**.
- Click **Refresh** to view your changes.

1. From the Upgrade section, click Choose File and select your file.
2. **HA Sync** - toggle **ON** to activate.
3. Click **Upload** . The upgrade process starts.

   After the new firmware has been installed, the system reboots.

When you update software, you are also updating the UI. To ensure the UI displays the updated pages correctly:

- Clear your browser cache.
- Refresh the page.

In most environments, press Ctrl-F5 to force the browser to get a new copy of the content from the web application. See the Wikipedia article on browser caching issues for a summary of tips for many environments: https://en.wikipedia.org/wiki/Wikipedia:Bypass_your_cache.

Deploy an active-passive cluster

In an active-passive cluster, one node is the active appliance; it processes traffic. The other node is passive; it is ready to assume the role of the active appliance if the primary node is unavailable.

Configure the system to send heartbeat packets between the pair to monitor availability. The system continually polls the activity of the heartbeat packets. If the active appliance becomes unresponsive, fail-over occurs: the standby becomes active. This process is shown below:

1. The standby node sends gratuitous ARP to notify adjacent routers to direct traffic for the virtual MAC addresses (vMAC) to its network interfaces.
2. It takes the IP addresses of the unresponsive node.

Figure 28 - An active-passive cluster at fail over - IP address transfer to the new active member

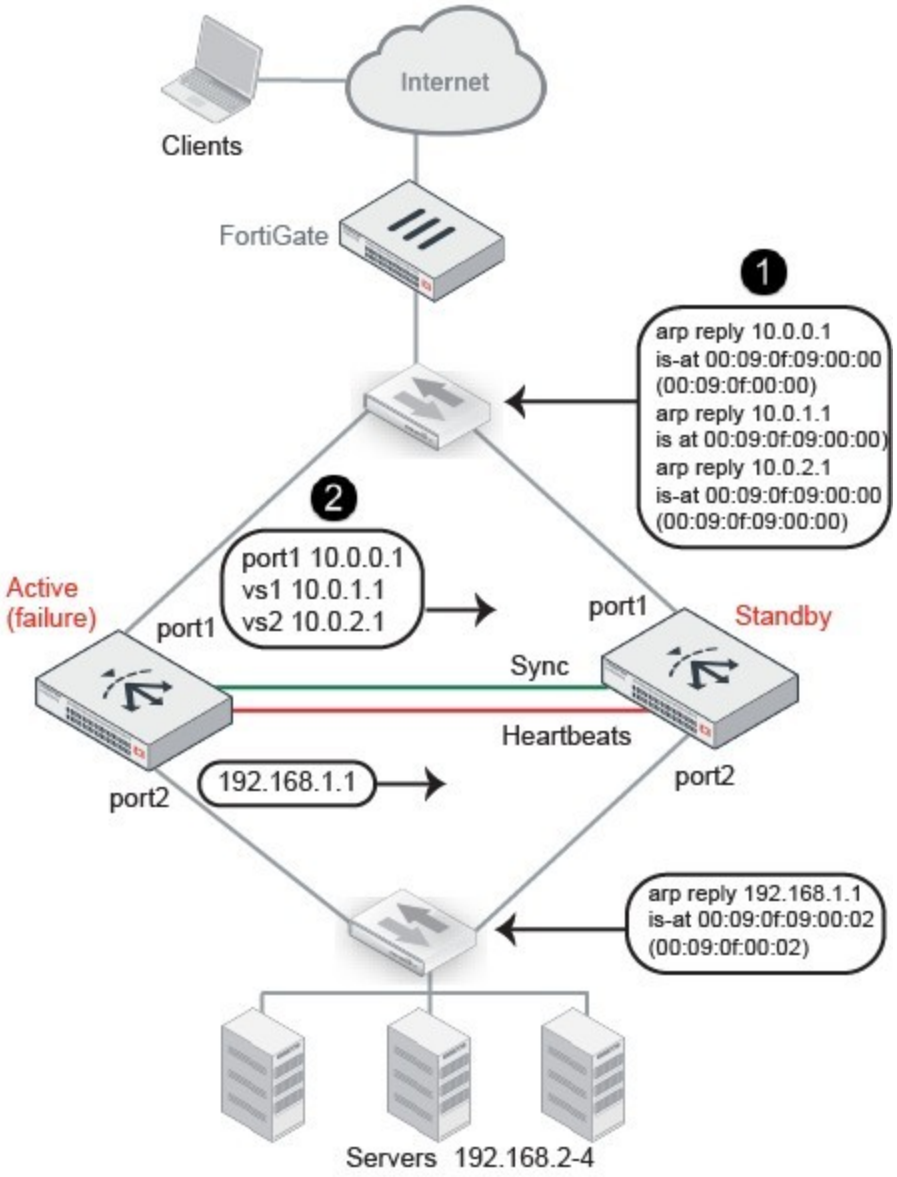

When the former active appliance comes back online, it might or might not assume its former active role. The system selects the active member based on the following criteria:

- Link health (if monitor ports links are down, the node is considered down)
- Remote IP monitor health check results
- Override setting (prefers priority to up time)
- Most available ports
- Highest up time value
- Lowest device priority number (1 has greater priority than 2)

- Highest-sorting serial number - Serial numbers are sorted by comparing each character from left to right, where 9 and z are the greatest values. The system gives preference to higher values over lower values.

Workflow

1. License all FortiWAN appliances in the HA cluster, and register them, including FortiGuard services, with the Fortinet Customer Service & Supportwebsite:

   https://support.fortinet.com/
2. Physically link the FortiWAN appliances that make up the HA cluster.

   You must link at least one of their ports (for example, port4 to port4) for heartbeat and synchronization traffic between members of the cluster. You can do either of the following:
3. Connect the two appliances directly with a crossover cable.
4. Link the appliances through a switch. If connected through a switch, the heartbeat interfaces must be reachable by Layer 2 multicast.
5. Configure the secondary node:
6. Log into the secondary appliance as the admin user.
7. Complete the HA settings as described in High Availability (HA) settings on page 45.

   Important: Set the Device Priority to a higher number than the preferred primary node; for example, set it to 2.
8. Configure the primary node:
   a. Log into the primary appliance as the admin user.
   b. Enter your information into the fields for all features, as well as the HA configuration.

      Important: Set the Device Priority to a lower number than the secondary node; for example, set it to 1.

After saving the HA configuration changes, cluster members join or rejoin the cluster. When configuration changes are saved on the primary node, its configuration is automatically pushed to the secondary node.

Active-Passive Cluster Best Practice

- Be careful to maintain the heartbeat links. If the heartbeat is accidentally interrupted, such as when a network cable is temporarily disconnected, the other nodes assume that the primary node has failed. In an active-passive deployment, fail over occurs. If no failure has actually occurred, both nodes can be operating as the active node. simultaneously.
- If you link HA appliances through switches, to improve fault tolerance and reliability, link the ports through two separate switches. Also, don't connect these switches to your overall

network, which could introduce a potential attack point, and could also allow network load to cause latency in the heartbeat, which could cause an unintentional fail over.

Log into a non-primary member node

In an Active-Passive cluster, only the management IP address for the primary node is active. In an active- passive cluster, you can log into a node only when it has primary node status and its IP address is active. To access the user interface of an appliance in standby status (the active-passive child), you must use a console port connection.

Use the `execute ha manage` command to log into the console of a member node from the master node.

Best Practices and Fine Tuning

For advice on best practices and performance, see these topics.

- Administrator access on page 30
- Regular backups on page 264
- The impact of logging on performance on page 196
- Packet capture on page 227
- Increase System performance on page 265
- Topology on page 267

About Network Security

In most deployment scenarios, we recommend you deploy FortiGate to secure your network. Fortinet includes security functionality within the FortiWAN system to support those times when deploying FortiGate is impractical. FortiWAN includes the following security features:

- **Firewall** - Drop traffic that matches a source-destination-application tuple you specify. See Firewall on page 176.
- **Connection limit** - Drop an abnormally high volume of traffic from a source-destination-application match. See Connection limit on page 172.

Regular backups

Make a backup before executing disruptive operations, such as:

- Upgrading the firmware.
- Running the CLI commands `execute factoryreset` or `execute restore`.
- Clicking the **Reset** button in the User Information widget on the dashboard.
- Always password-encrypt your backups.

Increase System performance

- Delete or disable unused policies. The system allocates memory with each server policy, regardless of whether it is actually in active use. Configuring extra policies will unnecessarily consume memory and decrease performance.
- To reduce latency associated with DNS queries, use a DNS server on your local network as your primary DNS.
- If the devices on your network support it, you can create one or more VLAN interfaces. VLANs reduce the size of a broadcast domain and the amount of broadcast traffic received by network hosts, thus improving network performance.

Troubleshooting

These topics offer troubleshooting tips.

Topology

FortiWAN uses Port1 and Port2, as the outbound interface, by default, is in DHCP mode. You can directly connect the network cable to connect the device to the Internet.

The Link Load balance module has a predefined **Link-Link group-Flow** policy that you can use without modification. See Link load balance (LLB) - Resources on page 113.

Make sure your FortiWAN device can connect to the Internet. By default, FortiWAN will actively detect 8.8.8.8 to determine network connectivity. If your FortiWAN device can’t connect to the Internet, we recommend that you disable the health check function in **Link**. See Health checks on page 92.

Login issues

If an administrator is entering his or her correct account name and password, but can't log in from some or all computers, examine the trusted host definitions for that account. It should include all locations where that person is allowed to log in, such as your office, but should not be too broad.

Connectivity issues

One of your first tests when configuring a new policy should be to determine whether allowed traffic is flowing to your servers. Investigate the following connectivity issues if traffic does not reach the destination servers:

- Is there a FortiWAN policy for the destination servers? By default, FortiWAN allows traffic to reach a backend server.

However, the virtual servers must also be configured before traffic can pass through.

- If your network utilizes secure connections (HTTPS) and there is no traffic flow, is there a problem with your certificate?

Check hardware connections

If there is no traffic flowing from the FortiWAN appliance, you want to rule out hardware problems.

- Ensure the network cables are properly plugged in to the interfaces on the FortiWAN appliance.
- Ensure there are connection lights for the network cables on the appliance.
- Change the cable if the cable or its connector are damaged or you are unsure about the cable's type or quality.
- Connect the FortiWAN appliance to different hardware to see if that makes a difference.
- In the UI, go to **Networking > Interface** and ensure the link status is up for the interface. If the status is down (down arrow on red circle), edit the configuration to change its status to Up.

You can also enable an interface using CLI commands. See CLI Configure Network Interfaces on page 237.

If any of these checks solve the problem, it was a hardware connection issue. You should still perform some basic software tests to ensure complete connectivity.

If the hardware connections are correct and the appliance is powered on but you can't connect using the CLI or UI, you may be experiencing bootup problems. See Restoring firmware ("clean install").

Check routing

The `ping` and `traceroute` utilities are useful for investigating issues with network connectivity

and routing.

Since you typically use these tools to troubleshoot, you can allow ICMP, the protocol used by these tools, in firewall policies and on interfaces only when you need them. Otherwise, disable ICMP for improved security and performance.

By default, FortiWAN appliances don't respond to `ping` and `traceroute`. However, if the appliance does not respond, and there are no firewall policies that block it, ICMP type 0 (`ECHO_REPSPONSE`) might be effectively disabled.

Enable ping and traceroute responses

1. Go to **Networking > Interface** and click **Edit** on the network interface row.
2. Under **Allow Access**, check **Ping**.
3. **Save**.

The appliance should now respond when another device such as your management computer sends a `ping` or `traceroute` to that network interface.

Disabling ping only prevents the system from receiving ICMP type 8 (`ECHO_ REQUEST`) and traceroute-related UDP. It does not disable CLI commands ,such as execute `ping` or execute `traceroute` that send such traffic.

Verify routes between clients and your servers

1. Try to connect through the FortiWAN appliance, from a client to a backend server, via HTTP or HTTPS. If the connectivity test fails, continue to the next step.
2. Use the `ping` command on both the client and the server to verify that a route exists between the two. Test traffic movement in both directions: from the client to the server, and the server to the client. Servers don't need to be able to initiate a connection, but must be able to send reply traffic along a return path.
   - If the routing test succeeds, check login or resource issues.
   - If the routing test fails, continue to the next step.
3. Use the `tracert` or `traceroute` command on both the client and the server (depending on their operating systems) to locate the point of failure along the route.

   If the route is broken when it reaches the FortiWAN appliance, first examine its network interfaces and routes. To display network interface addresses and subnets, enter the CLI command:

```
show system interface
```

You might need to verify that the physical cabling is reliable and not loose or broken, that there are no IP address or MAC address conflicts or blacklisting, misconfigured DNS records, and otherwise rule out problems at the physical, network, and transport layer.

If these tests succeed, a route exists, but you can't connect using HTTP or HTTPS, an application-layer problem is preventing connectivity.

4. For link load balance scenario, you might need to double check the LLB flow policies:

```
diagnose llb flow-policy list
```

5. For application-layer problems, on the FortiWAN, examine the:
   - virtual server policy and all components it references
   - server service/daemon

On routers and firewalls between the host and the FortiWAN appliance, verify that they permit HTTP and/or HTTPS connectivity between them.

Test connectivity with ping

The `ping` command sends a small data packet to the destination and waits for a response. The response has a timer that may expire, indicating that the destination is unreachable via ICMP.

ICMP is part of Layer 3 on the OSI Networking Model. ping sends Internet Control Message Protocol (ICMP) `ECHO_REQUEST` (“ping”) packets to the destination, and listens for `ECHO_RESPONSE` (“pong”) packets in reply.

Some networks block ICMP packets because they can be used in a `ping` flood or denial of service (DoS) attack if the network does not have anti-DoS capabilities, or because `ping` can be used by an attacker to find potential targets on the network.

Beyond basic existence of a possible route between the source and destination, pingtells you the amount of packet loss (if any), how long it takes the packet to make the round trip (latency), and the variation in that time from packet to packet (jitter).

If `ping` shows some packet loss, investigate:

- cabling to eliminate loose connections
- ECMP, split horizon, or network loops
- all equipment between the ICMP source and destination to minimize hops If `ping` shows total packet loss, investigate:
- cabling to eliminate incorrect connections
- all firewalls, routers, and other devices between the two locations to verify correct IP addresses, routes, MAC lists, and policy configurations

If `ping` finds an outage between two points, use `traceroute` to locate exactly where the problem is.

1. Log into the CLI via either SSH, Telnet, or the CLI Console widget of the UI.
2. To adjust the behavior of `execute ping`, first use the `execute ping-options` command.
3. Enter the command:

   ```
   execute ping <destination_ipv4>
   ```

   where `<destination_ipv4>` is the IP address of the device that you want to verify that the appliance can connect to, such as `192.168.1.1`.

If the appliance can reach the host via ICMP, output similar to the following appears:

```
PING 192.168.1.1 (192.168.1.1): 56 data bytes
64 bytes from 192.168.1.1: icmp_seq=0 ttl=253 time=6.5 ms
64 bytes from 192.168.1.1: icmp_seq=1 ttl=253 time=7.4 ms
64 bytes from 192.168.1.1: icmp_seq=2 ttl=253 time=6.0 ms
64 bytes from 192.168.1.1: icmp_seq=3 ttl=253 time=5.5 ms
64 bytes from 192.168.1.1: icmp_seq=4 ttl=253 time=7.3 ms

--- 192.168.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 5.5/6.5/7.4 ms
```

If the appliance can't reach the host via ICMP, output similar to the following appears:

```
PING 10.0.0.1 (10.0.0.1): 56 data bytes
Timeout ...
Timeout ...
Timeout ...
Timeout ...
Timeout ...

--- 10.0.0.1 ping statistics ---
5 packets transmitted, 0 packets received, 100% packet loss
```

“100% packet loss” and “Timeout” indicates that the host isn't reachable.

To verify that routing is bidirectionally symmetric, you should also ping the appliance.

Test routes and latency with traceroute

The traceroute utility sends ICMP packets to test each hop along the route. It sends three packets to the destination, and then increases the time to live (TTL) setting by one, and sends another three packets to the destination. As the TTL increases, packets go one hop farther along the route until they reach the destination.

Most traceroute commands display their maximum hop count—that's, the maximum number of steps it will take before declaring the destination unreachable—before they start tracing the route. The TTL setting may result in routers or firewalls along the route timing out due to high latency.

Where ping only tells you if the signal reached its destination and returned successfully, traceroute shows each step of its journey to its destination and how long each step takes. If you specify the destination using a domain name, the traceroute output can also indicate DNS problems, such as an inability to connect to a DNS server.

By default, the traceroute utility uses UDP with destination ports numbered from 33434 to 33534. The traceroute utility usually has an option to specify use of ICMP `ECHO_REQUEST` (type 8) instead, as used by the Windows tracert utility. If you have a firewall and you want traceroute to work from both machines (Unix- like systems and Windows) you will need to allow both protocols inbound through your firewall (UDP ports 33434 - 33534 and ICMP type 8).

1. Log into the CLI via either SSH, Telnet, or the CLI Console widget of the UI.
2. Enter the command:

   ```
   execute traceroute {<destination_ipv4> | <destination_fqdn>}
   ```

   where {`<destination_ipv4> | <destination_fqdn>`}is a choice of either the device's IP address or its fully qualified domain name (FQDN).

   For example, you might enter:

   ```
   execute traceroute www.example.com
   ```

If the appliance has a complete route to the destination, output similar to the following appears:

```
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84
bytepackets
1 172.16.1.2 0 ms 0 ms 0ms
2 209.87.254.221 <static-209-87-254-221.storm.ca> 2 ms 2 ms 2ms
3 209.87.239.129 <core-2-g0-1-1104.storm.ca> 2 ms 1 ms 2ms
4 67.69.228.161 2 ms 2 ms 3ms
5 64.230.164.17 <core2-ottawa23_POS13-1-0.net.bell.ca> 3 ms 3 ms 2ms
6 64.230.132.234 <core2-ottawatc_POS5-0-0.net.bell.ca> 20 ms 20 ms
20ms
7 64.230.132.58 <core4-toronto21_POS0-12-4-0.net.bell.ca> 24 ms 21
ms 24ms
```

```
8 64.230.138.154 <bx4-toronto63_so-2-0-0-0.net.bell.ca> 8 ms 9 ms
8ms
9 64.230.185.145 <bx2-ashburn_so2-0-0.net.bell.ca> 23 ms 23 ms 23ms
10 12.89.71.9 23 ms 22 ms 22ms
11 12.122.134.238 <cr2.wswdc.ip.att.net> 100 ms 12.123.10.130
<cr2.wswdc.ip.att.net> 101 ms 102ms
12 12.122.18.21 <cr1.cgcil.ip.att.net> 101 ms 100 ms 99ms
13 12.122.4.121 <cr1.sffca.ip.att.net> 100 ms 98 ms 100ms
14 12.122.1.118 <cr81.sj2ca.ip.att.net> 98 ms 98 ms 100ms
15 12.122.110.105 <gar2.sj2ca.ip.att.net> 96 ms 96 ms 96ms

16 12.116.52.42 94 ms 94 ms 94ms
17 203.78.181.10 88 ms 87 ms 87ms
18 203.78.181.130 90 ms 89 ms 90ms
19 66.171.121.34 <fortinet.com> 91 ms 89 ms 91ms
20 66.171.121.34 <fortinet.com> 91 ms 91 ms 89ms
```

Each line lists the routing hop number, the IP address and FQDN (if any) of that hop, and the 3 response times from that hop. Typically a value of `<1ms` indicates a local router.

If the appliance does not have a complete route to the destination, output similar to the following appears:

```
traceroute to 10.0.0.1 (10.0.0.1), 32 hops max, 84 byte packets
1 172.16.1.2 0 ms 0 ms 0ms
2 172.16.1.10 0 ms 0 ms 0ms
3 * * *
4 * * *
```

The asterisks (*) indicate no response from that hop in the network routing.

Examine the routing table

When a route does not exist, or when hops have high latency, examine the routing table. The routing table is where the FortiWAN appliance caches recently used routes.

If a route is cached in the routing table, it saves time and resources that would otherwise be required for a routelookup. If the routing table is full and a new route must be added, the oldest, least-used route is deleted to make room.

To check the routing table in the CLI, enter:

```
get router info routing-table all
```

Examine server daemons

If a route exists, but you can't connect to the UI using HTTP or HTTPS, an application-layer problem is preventing connectivity.

Verify that you have enabled HTTPS and/or HTTP on the network interface. Also examine routers and firewalls between the host and the FortiWAN appliance to verify that they permit HTTP and/or HTTPS connectivity between them. Finally, you can also use the CLI command to verify that the daemons for the UI and CLI, such as `sshd, cli, nginx`, and `php-fpm` are running and not overburdened:

```
diagnose system top delay 10
```

Checkport assignments

When attempting to connect to FortiWAN on a given network port, and the connection is expected to occur on a different port number, the attempt will fail. For a list of ports used by FortiWAN, see Port Numbers on page 284.

Performing a packet trace

When troubleshooting malformed packet or protocol errors, it helps to look inside the protocol headers of packets to determine if they are traveling along the route you expect, and with the flags and other options you expect.

If you configure virtual servers on your FortiWAN appliance, packets' destination IP addresses will be those IP addresses, not the physical IP addresses (i.e., the IP address of port1, etc.). An ARP update is sent out when a virtual IP address is configured.

If the packet trace shows that packets are arriving at your FortiWAN appliance's interfaces but no HTTP/HTTPS packets egress, check that:

- Physical links are firmly connected, with no loose wires
- Network interfaces are brought up
- Link aggregation peers, if any, are up
- VLANIDs, if any, match
- Virtual servers exist, and are enabled
- Matching policies exist, and are enabled
- If using HTTPS, valid server/CA certificates exist
- IP-layer and HTTP-layer routes, if necessary, match
- Servers are responsive, if server health checks are configured and enabled

Resource issues

Check these items in case of sluggish or stalled performance.

Monitoring traffic load

Heavy traffic loads can cause sustained high CPU or RAM usage. If this is unusual, no action is required. However, sustained heavy traffic load might indicate that you need a more powerful FortiWAN model.

In the UI, you can view traffic load two ways:

- Monitor current HTTP traffic on the dashboard. Go to **System > Dashboard > Link** and examine the throughput graphs.
- Examine traffic history in the traffic log. Go to **Logs > Log Browsing > Traffic Log**.

DoS attacks

A prolonged denial of service (DoS) can bring your servers down if your FortiWAN appliance and your network devices are not configured to prevent it. To prevent DoS attacks, enable the DoS and connection limit features. Also, configure protections on your FortiGate and other network devices. DoS attacks can use a variety of mechanisms. For in-depth protection against a wide variety of DoS attacks, you can use a specialized appliance such as FortiDDoS.

In the UI, you can watch for attacks in two ways:

- Monitor current traffic on the dashboard. Go to **System > Dashboard** and examine the system-wide throughput.
- Examine attack history in the traffic log. Go to **Log > Log Browsing > Security Log**.

Resetting the configuration

If you intend to sell your FortiWAN appliance, or if you are not sure what part of your configuration is causing a problem, you can reset it to its default settings and erase all data. If you have not updated the firmware, this is the same as resetting to the factory default settings.

Back up the configuration before performing a factory reset. See Back up and restore on page 42.

To **delete** your data from the system, connect to the CLI and enter this command:

```
execute formatlogdisk
```

To **reset** the configuration, connect to the CLI and enter this command:

```
execute factoryreset
```

Restoring firmware (“clean install”)

Restoring (also called re-imaging) the firmware can be useful if:

- you are unable to connect to the FortiWAN appliance using the UI or the CLI.
- you want to install firmware without preserving any existing configuration (i.e. a “clean install”).
- a firmware version that you want to install requires a different size of system partition (see the release notes accompanying the firmware).
- a firmware version that you want to install requires that you format the boot device (see the release notes accompanying the firmware).

This procedure applies to physical appliances. Restoring firmware re-images the boot device. Also, restoring firmware can only be done during a boot interrupt, before network connectivity is available, and therefore requires a local console connection to the CLI. It can't be done through an SSH or Telnet connection.

If you can't physically access the appliance’s local console connection, connect the appliance’s local console port to a terminal server to which you have net-work access. Once you have used a client to connect to the terminal server over the network, you are able to use the appliance local console through it. However, be aware that from a remote location, you might not be able to power cycle the appliance if abnormalities occur.

For virtual appliances, you can use VMware to backup and restore virtual appliance images.

Important: Back up the configuration before performing a clean install. See Back up and restore on page 42.

1. Download the firmware file from the Fortinet Customer Service & Support website: https://support.fortinet.com/.
2. Connect your management computer to the FortiWAN console port using a RJ-45-to-DB-9 serial cable or a null- modem cable.
3. Initiate a local console connection from your management computer to the CLI of the FortiWAN appliance, and log in as the admin administrator, or an administrator account whose access profile contains Read-Write permissions in the Maintenance category.
4. Connect port1 of the FortiWAN appliance directly or to the same subnet as a TFTP server.
5. Copy the new firmware image file to the root directory of the TFTP server.

6. If necessary, start your TFTP server. (If you don't have one, you can temporarily install and run one such as `tftpd` (Windows, Mac OS X, or Linux) on your management computer.)

TFTP isn't secure, and it does not support authentication. Only run it on trusted administrator-only networks, and never on computers directly connected to the Internet. Turn TFTP off immediately after completing this procedure.

7. Verify that the TFTP server is currently running, and that the FortiWAN appliance can reach the TFTP server. To use the FortiWAN CLI to verify connectivity, enter the following command:

```
execute ping 192.168.1.168
```

where `192.168.1.168` is the IP address of the TFTP server.

8. Enter the following command to restart the FortiWAN appliance:

```
execute reboot
```

As the FortiWAN appliances starts, a series of system startup messages appear.

```
Press any key to display configuration menu.....
```

9. Immediately press a key to interrupt the system startup.

You have only 3 seconds to press a key. If you don't press a key quickly enough, the FortiWAN appliance reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server. [F]:
Format boot device.
[B]: Boot with backup firmware and set as default.
[Q]: Quit menu and continue to boot with default firmware. [H]:
Display this list of options.
Enter G,F,B,Q,or H:
Please connect TFTP server to Ethernet port "1".
```

10. If the firmware version requires that you first format the boot device before installing firmware, type **F**.

Format the boot disk before continuing.

11. Type Gto get the firmware image from the TFTP server. The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

12. Type the IP address of the TFTP server and press **Enter**. The following message appears:

```
Enter local address [192.168.1.188]:
```

13. Type a temporary IP address that can be used by the FortiWAN appliance to connect to the TFTP server. The following message appears:

    Enter firmware image file name [image.out]:

14. Type the file name of the firmware image and press **Enter**.

    The FortiWAN appliance downloads the firmware image file from the TFTP server and displays a message similar to the following:

    ```
    MAC:00219B8F0D94
    ##########################
    Total 28385179 bytes data downloaded. Verifying the integrity of
    the firmware image..
    Save as Default firmware/Backup firmware/Run image without saving:
    [D/B/R]?
    ```

If the download fails after the integrity check with the error message:

```
invalid compressed format (err=1)
```

but the firmware matches the integrity checksum on the Fortinet CustomerService & Support website, try a different TFTP server.

15. Type **D**.

    The FortiWAN appliance downloads the firmware image file from the TFTP server. The FortiWAN appliance installs the firmware and restarts. The time required varies by the size of the file and the speed of your network connection.

    The FortiWAN appliance reverts the configuration to default values for that version of the firmware.

16. To verify that the firmware was successfully installed, log in to the CLI and type: `get system status`.

    The firmware version number shows.

    Either reconfigure the FortiWAN appliance or restore the configuration file.

Appendix

Fortinet MIBs

Listed are the management information bases (MIBs) used with FortiWAN.

MIB or RFC	Description
Fortinet Core MIB	This Fortinet-proprietary MIB enables your SNMP manager to query for system information and to receive traps that are common to multiple Fortinet devices.
FortiWAN	This Fortinet-proprietary MIB enables your SNMP manager to query for FortiWAN-specific information and to receive FortiWAN-specific traps.
RFC 1213 (MIB II)	The FortiWAN SNMP agent supports MIB II groups, except: There is no support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). Protocol stat- istics returned for MIB II groups (IP, ICMP, TCP, UDP, and so on) don't accur- ately capture all FortiWAN traffic activity. More accurate information can be obtained from the information reported by the FortiWAN MIB.
RFC 3635 (Ethernet-like MIB)	The FortiWAN SNMP agent uses any of the objects in the Ethernet-like interface

Download the Fortinet MIB files from the Fortinet Customer Service & Support website, https://support.fortinet.com/.

To view a trap or query's name, object identifier (OID), and description, open its MIB file in a plain text editor.

To communicate with the FortiWAN SNMP agent, you must first compile these MIBs into your SNMP manager. If the standard MIBs used by the SNMP agent are already compiled into your SNMP manager, you don't have to compile them again. The FortiWAN SNMP implementation is read-only.

All traps sent include the message, the FortiWAN appliance's serial number, and host name.

Figure 29 - FortiWAN MIB download

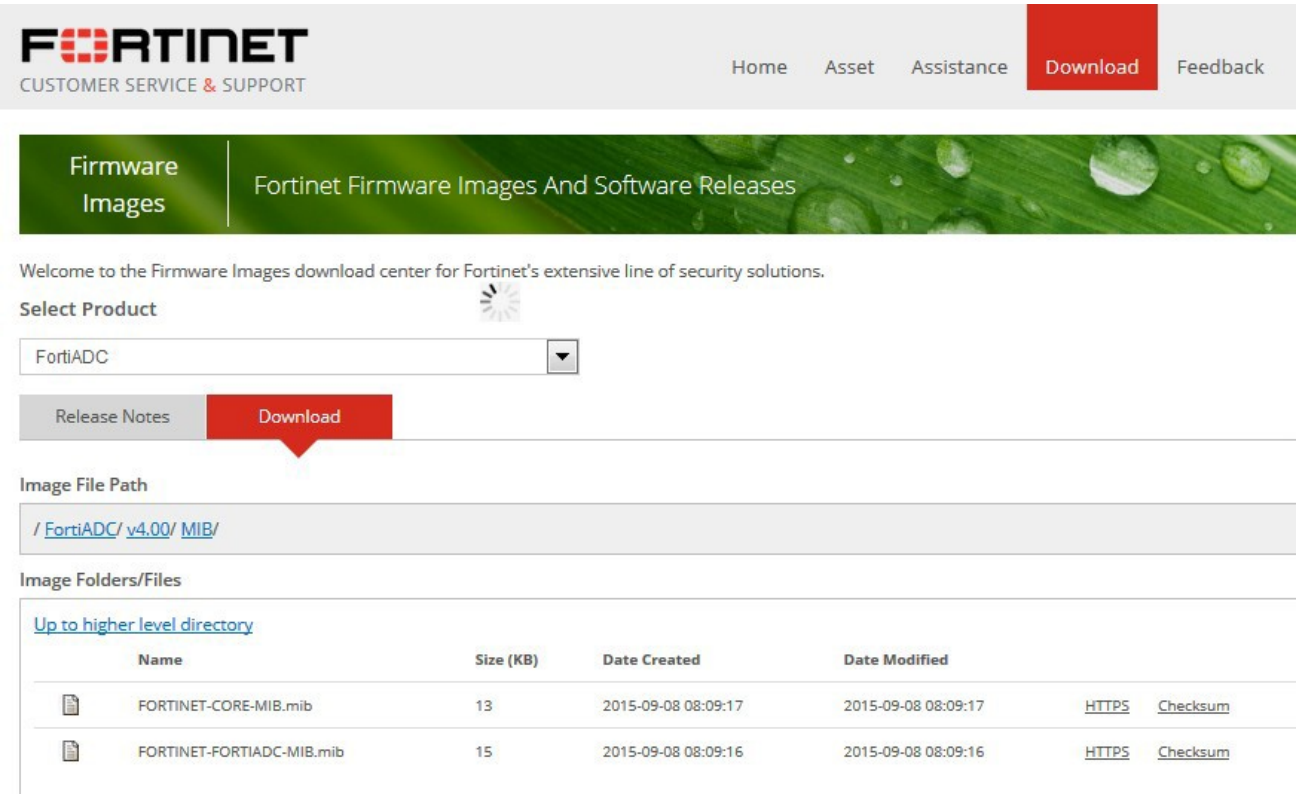

Port Numbers

Communications between the FortiWAN system, clients, servers, and FortiGuard Distribution Network (FDN) require that any routers and firewalls between them permit specific protocols and port numbers.

The following tables list the default port assignments used by the FortiWAN system.

Default ports used by FortiWAN for outgoing traffic

Port Number	Protocol	Purpose
N/A	ARP	HA fail over of network interfaces.
N/A	ICMP	Server health checks. `execute ping` and `execute traceroute`.
25	TCP	SMTP for alert email.
53	UDP	DNS queries.
69	UDP	TFTP for backups, restoration, and firmware updates. See commands such as execute back up or execute restore.
80	TCP	Server health checks.
123	UDP	NTP synchronization.
162	UDP	SNMP traps.
389	TCP	LDAP authentication queries.
443	TCP	FortiGuard polling. Server health checks.
514	UDP	Syslog.
6055	UDP	HA heartbeat. Layer 2 multicast.

Port Number	Protocol	Purpose
6056	UDP	HA configuration synchronization. Layer 2 multicast.

Default ports used by FortiWAN for incoming traffic (listening)

Port Number	Protocol	Purpose
N/A	ICMP	ping and traceroute responses.
22	TCP	SSH administrative CLI access.
23	TCP	Telnet administrative CLI access.
53	UDP	DNS queries from clients for global load balancing and inbound link load balancing.
80	TCP	HTTP administrative UI access. Predefined HTTP service. Only occurs if the service is used by a virtual server.
161	UDP	SNMP queries.
443	TCP	• **HTTPS administrative web UI access**. Only occurs when the destination address is a network interface IP address. • **Predefined HTTPS service**. Only occurs if the service is used

Port Number	Protocol	Purpose
		by a virtual server, and if the destination address is a virtual server.
6055	UDP	HA heartbeat. Layer 2 multicast.
6056	UDP	HA configuration synchronization. Layer 2 multicast.

Maximum Configuration Values

This table shows the maximum number of configuration objects or limits, and are not a guarantee of performance. For values such as hardware specifications that don't vary by software version or configuration, see the Quick Start Guide or data sheet for your model.

Maximum configuration objects - Hardware models

		30E	200E
System			
Administration	Administrative users	300	300
	Access profiles	8	16
Shared Resources	Address	4096	4096
	Address group	4096	4096
	Health checks	128	128
	ISP address book	32	32
	Schedule	256	256
	Schedule group	64	64
	Application	4096	4096
	Application group	1024	1024
SNMP	SNMP community	16	16
	SNMP community Host	16	16
	SNMP user	16	16
Networking			
Interface	Physical network interfaces	5	6

		30E	200E
	Total interfaces	256	1024
IPSEC	Ipsec phase1	128	512
DHCP	DHCP server	32	64
ARP	ARP table entries	4096	4096
Routing	Static routes	512	1024
	Policy routes	32	64
	Any configuration object	256	256
Link Load Balancing			
	Link	256	512
	Link group	256	512
	Link group member	256	512
	Flow policy rule	512	1024
Firewall			
	Firewall policy	256	512
Connection Limit			
	Connection limit policy	256	512
NAT			
	Snat	1024	2048
	Vip	32	64
Bandwidth Management			
	Traffic shaper	512	1024
	Traffic shaper policy	512	1024

		30E	200E
Server Load Balancing			
Virtual Servers	Virtual Server	32	64
Real Server	Real server	32	64
	Server pool	32	64
	Pool members	32	64
Resources	Profiles	32	32
	Persistence	32	32
	IP pool	16	32
Global DNS Server			
	Any configuration object	256	256
Log & Report			
	Remote Syslog Servers	3	3
	IPFIX Servers	3	3
Diagnostics			
Packet capture		5	5

Maximum configuration objects - Virtual Appliances

		VM10/20	VM50/1H	VM2H/5H	VM1K/2K
System					
Administration	Administrative users	300	300	300	300
	Access profiles	8	16	64	64

		VM10/20	VM50/1H	VM2H/5H	VM1K/2K
Shared Resources	Address	4096	4096	4096	4096
	Address group	4096	4096	4096	4096
	Health checks	128	128	128	128
	ISP address book	32	32	32	32
	Schedule	256	256	256	256
	Schedule group	64	64	64	64
	Application	4096	4096	4096	4096
	Application group	1024	1024	1024	1024
SNMP	SNMP community	16	16	16	16
	SNMP community Host	16	16	16	16
	SNMP user	16	16	16	16
Networking					
Interface	Physical network interfaces	10	10	10	10
	Total interfaces	256	1024	4096	8192
IPSEC	Ipsec phase1	128	512	2048	4096
DHCP	DHCP server	32	64	128	256
ARP	ARP table entries	4096	4096	4096	4096
Routing	Static routes	512	1024	2048	4096
	Policy routes	32	64	128	256
	Any configuration object	256	256	256	256
Link Load Balancing					

		VM10/20	VM50/1H	VM2H/5H	VM1K/2K
	Link	256	512	1024	2048
	Link group	256	512	1024	2048
	Link group member	256	512	1024	2048
Firewall					
	Firewall policy	256	512	1024	2048
Connection Limit					
	Connection limit policy	256	512	1024	2048
NAT					
	Snat	1024	2048	4096	8192
	Vip	32	64	128	256
Bandwidth Management					
	Traffic shaper	512	1024	2048	4096
	Traffic shaper policy	512	1024	2048	4096
Server Load Balancing					
Virtual Servers	Virtual Server	32	64	128	256
Real Server	Real server	32	64	128	256
	Server pool	32	64	128	256
	Pool members	32	64	128	256
Resources	Profiles	32	32	32	32
	Persistence	32	32	32	32
	IP pool	16	32	64	128
Global DNS Server					
	Any configuration	256	256	256	256

		VM10/20	VM50/1H	VM2H/5H	VM1K/2K
	object				
Log & Report					
	Remote Syslog Servers	3	3	3	3
	IPFIX Servers	3	3	3	3
Diagnostics					
	Packet capture	5	5	5	5

Additional resources

For more information, refer to:

- The **Release Notes** provided with your firmware
- Technical documentation (reference guides, installation guides, and other documents)
- Knowledge base (technical support articles)

Forums

- Online campus (tutorials and training materials)

If you have problem using FortiWAN, check within your organization first. You can save time and effort during the troubleshooting process by checking if other FortiWAN administrators have experienced a similar problem before.

If you can't resolve the issue on your own, contact Fortinet Customer Service & Support.

Index

R

S

T

U

V

W

List of Figures