



---

# vGW Series High Availability and Fault Tolerance



---

Published: 2015-02-19

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Copyright © 2015, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*vGW Series High Availability and Fault Tolerance*  
Copyright © 2015, Juniper Networks, Inc.  
All rights reserved.

The information in this document is current as of the date on the title page.

#### YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

#### END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Table of Contents

	About the Documentation . . . . .	ix
	Documentation and Release Notes . . . . .	ix
	Documentation Conventions . . . . .	ix
	Documentation Feedback . . . . .	xi
	Requesting Technical Support . . . . .	xii
	Self-Help Online Tools and Resources . . . . .	xii
	Opening a Case with JTAC . . . . .	xii
<b>Part 1</b>	<b>High Availability</b>	
<b>Chapter 1</b>	<b>vGW Series Solution . . . . .</b>	<b>3</b>
	Understanding the vGW Series High Availability Solution . . . . .	3
	vGW Series HA . . . . .	3
	vGW Series HA and VMware HA . . . . .	3
	vGW Series HA for the vGW Security Design VM . . . . .	4
	vGW Security Design VM HA Behavior . . . . .	6
	vGW Series HA for the vGW Security VM . . . . .	8
<b>Chapter 2</b>	<b>Secondary Components Installation . . . . .</b>	<b>11</b>
	Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability . . . . .	11
	Installing a Secondary vGW Security VM for High Availability . . . . .	14
<b>Part 2</b>	<b>Fault Tolerance</b>	
<b>Chapter 3</b>	<b>vGW Series Handling of VMs Enabled for FT . . . . .</b>	<b>19</b>
	Understanding vGW Series Fault Tolerance Support . . . . .	19
	About vGW Series Fault-Tolerance . . . . .	19
	vGW Series Fault Tolerance in the vGW Series . . . . .	20
	Enabling Fault Tolerance for a Virtual Machine . . . . .	20
<b>Part 3</b>	<b>Index</b>	
	Index . . . . .	25



# List of Figures

<b>Part 1</b>	<b>High Availability</b>	
<b>Chapter 1</b>	<b>vGW Series Solution</b> .....	<b>3</b>
	Figure 1: Configuring Network Monitoring and NetFlow Settings .....	8
<b>Chapter 2</b>	<b>Secondary Components Installation</b> .....	<b>11</b>
	Figure 2: Configuring the Secondary vGW Security Design VM .....	12



# List of Tables

<b>About the Documentation</b> .....	<b>ix</b>
Table 1: Notice Icons .....	x
Table 2: Text and Syntax Conventions .....	x





# About the Documentation

- Documentation and Release Notes on page ix
- Documentation Conventions on page ix
- Documentation Feedback on page xi
- Requesting Technical Support on page xii

## Documentation and Release Notes

---

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

## Documentation Conventions

---

Table 1 on page x defines notice icons used in this guide.

Table 1: Notice Icons







Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page x defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"><li>To configure a stub area, include the <b>stub</b> statement at the <b>[edit protocols ospf area area-id]</b> hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options {   static {     route default {       nexthop <i>address</i>;       retain;     }   } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>.

- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## Requesting Technical Support

---

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.



## PART 1

# High Availability

- [vGW Series Solution on page 3](#)
- [Secondary Components Installation on page 11](#)





## CHAPTER 1

# vGW Series Solution

- [Understanding the vGW Series High Availability Solution on page 3](#)

## Understanding the vGW Series High Availability Solution

---

This topic gives an overview of the vGW Series high availability (HA) feature. It includes the following sections:

- [vGW Series HA on page 3](#)
- [vGW Series HA and VMware HA on page 3](#)
- [vGW Series HA for the vGW Security Design VM on page 4](#)
- [vGW Security Design VM HA Behavior on page 6](#)
- [vGW Series HA for the vGW Security VM on page 8](#)

## vGW Series HA

vGW Series provides high availability support for VMware environments for both the vGW Security Design VM and vGW Security VMs. The high availability feature maintains solution resiliency in the event of a failure. It allows you to deploy primary and secondary, or standby, vGW Security Design VMs and vGW Security VMs in which the secondary instance of the component takes control if the primary one is unavailable. vGW Series HA is effective in situations in which both primary components are inactive or only one is.



**NOTE:** vGW Series HA is an optional component which requires a separate license. You must purchase a separate 'VGW-HA' license for each vGW Security VM for which you plan to use the feature. The license allows the use of vGW HA for both the vGW Security Design VM and the vGW Security VM. You do not need to buy additional licenses for vGW Security VMs.

## vGW Series HA and VMware HA

vGW Series is compatible with VMware HA. You can configure any regular VM in your virtualized environment with VMware HA and still protect it with vGW Series security. Additionally, you can configure the vGW Security Design VM for VMware HA or fault tolerance (FT). When it is in effect, the VMware vCenter heartbeat does not impact vGW Series adversely.



**NOTE:** It is neither necessary nor possible to configure VMware HA or FT on vGW Security VMs.

vGW Series HA maintains two separate vGW Security VMs. It checks the health between these systems. If for some reason an OS or service crash occurs in the primary vGW Security VM, the secondary vGW Security VM takes over functionality.

### vGW Series HA for the vGW Security Design VM

The vGW Security Design VM, also referred to as the management center, is the main point of control for the entire vGW Series infrastructure. It presents the vGW Series interface to users, and it implements firewall security by distributing policy to the vGW Security VMs that protect ESX/ESXi hosts. You use it to configure the features that vGW Series provides and to view the wide range of information reported in its graphs, charts, and statistics. It consolidates logging information and it hosts the network monitoring database. If the vGW Security Design VM is unavailable, for example, because it crashed or it was turned off, an administrator cannot make configuration changes to the infrastructure nor benefit from information that the vGW Security Design VM gathers from virtualized environment and reports on. To protect against your inability to access this information, you can configure vGW Series HA support to enable a secondary vGW Security Design VM to take over when the primary one is unavailable.

vGW Series option to deploy both primary and secondary vGW Security Design VMs allows the secondary vGW Security Design VM to continue to serve up policy until the primary one can be brought back online. As a result, all normal network activity can continue without interruption, and new VMs powered on ESX/ESXi hosts can retrieve policy rather than defaulting to VMware failure mode.



**NOTE:** vGW Series high availability is meant to be used as an emergency solution, not as a replacement system. If the primary vGW Security Design VM fails, it can be recovered from a backup or snapshot copy. For details, see *Configuring the vGW Series Backup and Restore Feature*.

After you use the Settings module vGW Application Settings > High Availability page to select the vGW Security Design VM to use as the secondary one, the secondary vGW Security Design VM is automatically powered on and configured. The process takes approximately ten minutes.

“Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability” on page 11 explains the process for creating a secondary vGW Security Design VM.

The standby vGW Security Design VM presents the same address configuration options. Supported address types include:

- IPv4:

For IPv4, from the displayed list, select the method to use to assign an IPv4 address to Interface 1:

- **DHCP**

Use a DHCP server to assign dynamically an IPv4 address to Interface 1. This is the default method.

- **Static IP**

Specify a static IP address and its network mask routing prefix, and the default gateway to assign to Interface 1.

- **IPv6:**

For IPv6, from the displayed list, select the method to use to assign an IPv6 address to Interface 1:

- **DHCPv6**

Use a DHCPv6 server to obtain the IPv6 address for Interface 1. This is the default method.

Dynamic Host Configuration Protocol for IPv6 (DHCPv6) offers the capability of automatic allocation of reusable network addresses and additional configuration flexibility. This protocol is a stateful counterpart to IPv6 stateless address autoconfiguration.

- **Autoconfiguration**

Use stateless address autoconfiguration to obtain the IPv6 address for Interface 1. IPv6 stateless address autoconfiguration allows network devices attached to an IPv6 network to automatically acquire IP addresses and connect to the Internet without intermediate interaction with a DHCPv6 server.

Refer to *RFC 2462, IPv6 Stateless Address Autoconfiguration* for details.

- **Static IP**

Specify a static IP address for Interface 1 including the IPv6 address prefix (the initial bits of the address that denote the network address, akin to a netmask) and the default gateway to use for it.

When you configure the address for the secondary vGW Security Design VM, you must use the address type that you used to configure the primary vGW Security Design VM. However, if, for some reason, the address type configuration differs, you need to take into consideration problems that can ensue.

In an environment in which the vGW Security Design VM is configured for dual stack communication and you configure the secondary, or standby, vGW Security Design VM differently, that is, not for dual stack, communication problems should not occur. However, problems will occur if both the primary vGW Security Design VM and the standby vGW Security Design VM are not configured for dual stack and the protocol types of the IP addresses bound to them differ.

When your environment has a standby vGW Security Design VM that has only an IPv6 address bound to it, if you attempt to change the primary vGW Security Design VM from dual stack to single with only an IPv4 address bound to it, vGW Series displays the following message:

"The interface for management communications must have an IPv6 configuration, because there is a Standby Appliance with IPv6 interface."

See ["Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability"](#) on page 11.



**NOTE:** By default, a dual stack vGW Security Design VM communicates with a vGW Security VM using the IPv4 protocol. However, you can use the vGW CLI to change the default IP protocol used by setting the `center.dual.stack.default.communication.ipv4` parameter to false.

**`center.dual.stack.default.communication.ipv4=false`**

By default this parameter is set to true.

This parameter is relevant only if the vGW Security Design VM is configured for dual stack and one or more vGW Security VMs is also configured for dual stack. In all other cases, the protocol used is the one that is common to both the vGW Security Design VM and the vGW Security VM, and this parameter is irrelevant.

---

## vGW Security Design VM HA Behavior

vGW Series high availability for the vGW Security Design VM behaves in the following ways:

- It allows the secondary vGW Security Design VM to continue to distribute policy until the primary one can be brought back online. In this case, the term policy is used in a broad sense; it is meant to include vGW AntiVirus policy as well as firewall policy. When the primary vGW Security Design VM is unavailable, the secondary vGW Security Design VM pushes out the policy database to the vGW Security VMs when they request it. This policy is a copy of what existed in the primary vGW Security Design VM. It cannot be modified.

You cannot view anything related to compliance rules, network monitoring statistics, IDS, or vGW AntiVirus.

The high availability capability is intended to be used for emergency situations to ensure that new VMs that are powered on ESX/ESXi hosts can retrieve policy rather than default to VMsafe failure mode.



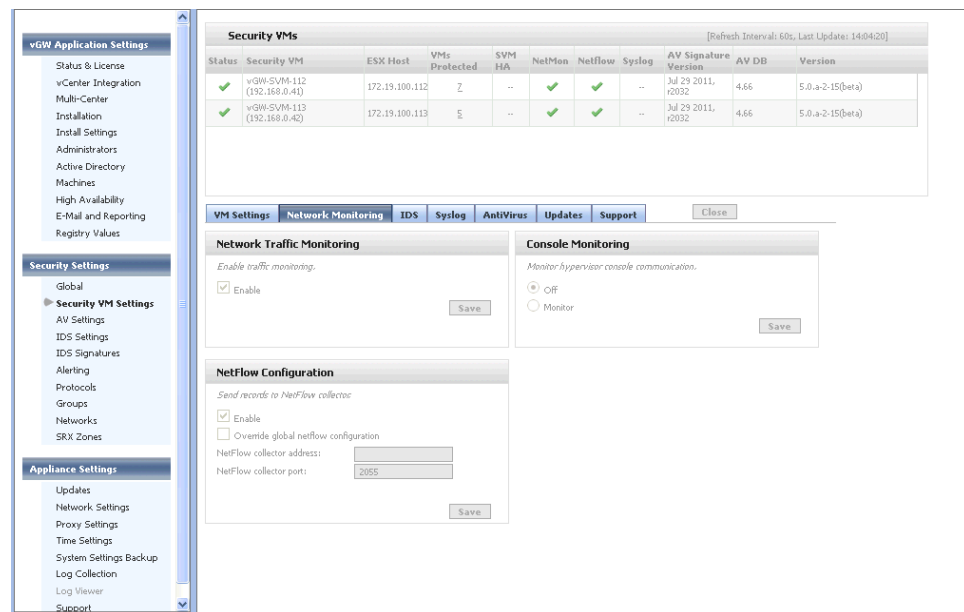
**NOTE:** It is important to understand that you cannot control features from the secondary vGW Security Design VM in ways in which you can using the primary one. You cannot configure new policies or modify existing ones.

- vGW Series does not synchronize events back from the secondary vGW Security Design VM to the primary one.
  - You cannot create compliance rules.
  - Changes to Network Monitoring and NetFlow, IDS, and AntiVirus events and changes to statistics are not viewable unless you configure NetFlow and Syslog from the vGW Security VM for individual vGW Security VMs. Although you cannot view their activity from the secondary vGW Security Design VM, these features continue to work when the primary vGW Security Design VM is unavailable.

If Network Monitoring is not enabled for the primary vGW Security VM, Console Monitoring is turned off. Network Monitoring must be enabled for Console Monitoring to work. These features continue to work as configured for the vGW Security VM when the primary vGW Security Design VM is unavailable.

[Figure 1 on page 8](#) shows the Settings module page that you use to configure Network Monitoring and NetFlow for individual vGW Security VMs.

Figure 1: Configuring Network Monitoring and NetFlow Settings



- Compliance and Introspection tasks, which rely on the primary vGW Security Design VM, are inactive.
- Updates to IDS signatures are not made.
- Updates to AntiVirus continue to occur.

## vGW Series HA for the vGW Security VM

In addition to providing for a secondary vGW Security Design VM, it is important to have redundancy at the vGW Security VM level. A vGW Security VM might become inactive, for example, when the vGW Security Design VM is inactive and its secondary takes over.

When the primary vGW Security VM becomes inactive, the secondary one becomes active in 60 seconds.

High availability considerations for the vGW Security VM differ from those of the vGW Security Design VM.

The secondary vGW Security VM is the same as the primary one, and it has the same capability, given certain circumstances.

- If the primary vGW Security Design VM is active and high availability is configured for a vGW Security VM, when a primary vGW Security VM becomes inactive other vGW Security VMs can perform introspection scans on behalf of its secondary.

It is also possible for the primary vGW Security Design VM to participate in the process, if it is active. (The secondary vGW Security Design VM cannot do this.)

- AntiVirus remains in effect, and AntiVirus signature updates take place regardless of whether the primary vGW Security Design VM is active.

A vGW Security VM is installed on each ESX/ESXi host to be protected. It is designed to interface directly with the hypervisor on its host. It is responsible for protecting VMs only on its host. Because of the tight coupling of a vGW Security VM and its host, it is important that a vGW Security VM not be moved to a new ESX/ESXi host. If the host is down, there is nothing to be protected.

Problems can occur if a vGW Security VM is not reinstated to its original position after failure. To protect against potential problems in this area, the vGW Series automatically sets the VMware high availability and Distributed Resource Schedule (DRS) settings to restrict vGW Security VMs from being moved through high availability or DRS.

To install a secondary vGW Security VM, you build another virtual machine from the original vGW Security VM. Unlike the process for creating a secondary vGW Security Design VM anew, when you create a secondary vGW Security VM, vGW Series clones the existing vGW Security VM.

For details on how to install a vGW Security VM, see [“Installing a Secondary vGW Security VM for High Availability” on page 14](#).

It is important to consider that the IP protocol address type of the IP address bound to the management interface of the secondary vGW Security VM must correspond to that of the vGW Security Design VM management interface with which it communicates. However, if both or either one is configured for dual stack, communication problems should not occur. If both are not configured for dual stack and the types of the IP addresses bound to their management interfaces differs, communication problems will ensue. For further information, see *Installing vGW Security VMs on ESX/ESXi Hosts*.

This pane allows you to change the IP protocol family that is used for the vGW Security VM management interface when that protocol does not match that of the vGW Security Design VM with which it must communicate. For information on conditions that would cause an IP address type mismatch between the management interfaces of the vGW Security VM and the vGW Security Design VM, see *Setting Up vGW Series* and *Installing vGW Security VMs on ESX/ESXi Hosts*.

**Related  
Documentation**

- *Understanding vGW Series*
- *Understanding the vGW Security Design VM*
- *Understanding the vGW Security VM*
- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 11](#)
- [Installing a Secondary vGW Security VM for High Availability on page 14](#)
- *Adding and Editing vGW Series Machines Definitions (VMware)*





## CHAPTER 2

# Secondary Components Installation

- Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 11
- Installing a Secondary vGW Security VM for High Availability on page 14

## Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability

---

This topic explains how to install an additional vGW Security Design VM to be used when the primary one is unavailable. You can install more than one additional vGW Security Design VM. It also explains how to configure the primary vGW Security Design VM for HA and how to determine the secondary one to use for it. The process entails:

- building another vGW Security Design VM from the vGW Series OVA file.
- selecting and configuring the secondary vGW Security Design VM to use for the primary one on the Settings module vGW Application Settings > High Availability page.



**CAUTION:** Be sure to back up your primary vGW Security Design VM. vGW Series does not rebuild a primary vGW Security Design VM from a secondary one created for HA. For details on backing up the primary vGW Security Design VM, see *Configuring the vGW Series Backup and Restore Feature*.

To create a secondary vGW Security Design VM:

1. Load the OVA file for the vGW Security Design VM using the VMware vSphere Client. (Use **File > Virtual Appliance > Import** in VMware vCenter.)
2. Follow the Virtual Appliance Wizard process. Accept the defaults for the virtual appliance import.



**NOTE:** If you need further information about installing the secondary vGW Security Design VM, you can read about how it is done for the primary vGW Security Design VM. See *Understanding the Open Virtualization Format OVA Template Method* and *Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware*, and related topics that they refer to.

The OVA import process prompts you for a database disk. You can accept the default 8.0 GB size even if your primary vGW Security Design VM is configured for a larger size. The secondary vGW Security Design VM does not store the same type of information as the primary one. Therefore it does not require more than 8.0 GB capacity.



**CAUTION:** After the import completes, do not power on the newly created secondary vGW Security Design VM.

To configure the primary vGW Security Design VM for HA:

1. Configure the vGW Security Design VM for HA in the Settings module:
  - a. To configure the secondary vGW Security Design VM, select **vGW Application Settings > High Availability**. See [Figure 2 on page 12](#).

**Figure 2: Configuring the Secondary vGW Security Design VM**

The screenshot shows the Juniper vGW Security Design VM configuration interface. The left sidebar has a menu with 'vGW Application Settings' expanded, showing options like 'Status & License', 'vCenter Integration', 'Multi-Center', 'Installation', 'Install Settings', 'Administrators', 'Active Directory', 'Machines', 'High Availability' (selected), 'E-Mail and Reporting', and 'Registry Values'. Below this is 'Security Settings' with options like 'Global', 'Security VM Settings', 'AV Settings', 'IDS Settings', 'IDS Signatures', and 'Alerting'. The main panel is titled 'Custom Standby Appliance' and contains a 'Standby Appliance' dropdown menu. Below it is the 'Standby Appliance management interface' section with fields for 'Internet Protocol' (set to DHCP), 'IP Address', 'Netmask', 'Default Gateway', and 'IPv6' (set to Disabled). There is a 'Save' button at the bottom right of this section. Below the interface section is a 'Settings' section with two checkboxes: 'Use Proxy Settings from Primary on Standby Center' and 'Use Time configuration settings from Primary on Standby Center', both of which are checked. There is a 'Save' button at the bottom right of the Settings section.

- b. From the Standby Appliance list, select the vGW Security Design VM to be used as the secondary (standby) vGW Security Design VM.
- c. Select the IP address type to assign to the secondary vGW Security Design VM and how it will obtain the address. You can select an IPv4 or IPv6 address.



**NOTE:** IPv4 DHCP is enabled by default.

From the Internet Protocol list select:

- For IPv4
  - **Disabled**

Disable IPv4 and use an IPv6 address for the secondary vGW Security Design VM.

- **DHCP**

Use DHCP to assign an IPv4 address dynamically to the secondary vGW Security Design VM.

- **Static IP**

Specify a static IPv4 address, its network mask routing prefix, and the default gateway to use for the secondary vGW Security Design VM.

- For IPv6:

- **Disabled**

Disable IPv6 and assign an IPv4 address to the secondary vGW Security Design VM.

- **DHCPv6**

Use a DHCPv6 server to obtain the IPv6 address to assign to the secondary vGW Security Design VM.

According to RFC 3315, "The Dynamic Host Configuration Protocol for IPv6 (DHCP) enables DHCP servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes. It offers the capability of automatic allocation of reusable network addresses and additional configuration flexibility. This protocol is a stateful counterpart to "IPv6 Stateless Address Autoconfiguration" (RFC 2462), and can be used separately or concurrently with the latter to obtain configuration parameters."

- **Autoconfiguration**

Use stateless address autoconfiguration to obtain the IPv6 address for the secondary vGW Security Design VM. IPv6 stateless address autoconfiguration allows network devices attached to an IPv6 network to automatically acquire IP addresses and connect to the Internet without intermediate interaction with a DHCPv6 server. Refer to RFC 2462, "IPv6 Stateless Address Autoconfiguration" for details.

- **Static IP**

Specify a static IPv6 address, its prefix (the initial bits of the address that denote the network address, akin to a netmask), and the default gateway to use for the secondary vGW Security Design VM.

d. Click **Save**.

2. Configure the proxy server and time configuration settings for the secondary vGW Security Design VM.

a. Specify whether to use the proxy settings configured for the primary vGW Security Design VM for the standby (secondary) one. See *Configuring vGW Series Proxy Settings*.

The Security Design vGW connects to the Juniper Networks update server to check for available downloads of software updates. If the server does not have direct access to the Internet, a proxy can be used. For the primary vGW Security Design

VM, the Settings module Appliance Settings > Proxy Settings page specifies configuration information about a proxy server, if one is required to make outbound http/https connections.

- b. Specify whether to use the time configuration settings configured for the primary vGW Security Design VM on the standby (secondary) one. See *Configuring vGW Series Time Settings*.
- c. Click **Save**.

After you complete this configuration, the secondary vGW Security Design VM is automatically powered on and configured. This process takes approximately ten minutes. After the operation completes, you can log in to the secondary vGW Security Design VM through the IP address that you specified during the configuration.

vGW Series monitors connectivity between the two vGW Security Design VM management centers. It initiates promotion of the secondary system if there is no response from the primary one within three minutes.

When the primary vGW Security Design VM is brought back online after it has recovered or the host it was on is repaired, it automatically takes control again. vGW Series HA is not designed to replace normal backup operations. Rather, it is expected that the primary vGW Security Design VM will be brought back online quickly.

**Related  
Documentation**

- [Understanding the vGW Series High Availability Solution on page 3](#)
- [Understanding the vGW Security Design VM](#)
- [Installing a Secondary vGW Security VM for High Availability on page 14.](#)
- [Preparing to Integrate the vGW Series with the VMware Environment](#)
- [Understanding vGW Series Fault Tolerance Support on page 19](#)

---

## Installing a Secondary vGW Security VM for High Availability

To install a secondary vGW Security VM for high availability (HA), you build another one from the original vGW Security VM. vGW Series clones the original vGW Security VM to create the standby one for HA.

HA for the vGW Security VM differs from HA for the vGW Security Design VM in the following ways:

- When you create a new vGW Security VM, vGW Series clones the existing one. You do not need to install a second template to generate it.
- It is not important to back up the vGW Security Design VM. If it is necessary, you can create another one from the template used to generate the original vGW Security VM.

Though the method of recreating the vGW Security Design VM, vGW Series does not rebuild the original vGW Security VM from the secondary VM.

However, like vGW Security Design VM, vGW Series does not rebuild the original vGW Security VM from the secondary VM.

To clone the existing, primary vGW Security VM:

1. Select Settings > Security Settings > Security VM Settings.
2. Click the row for the vGW Security VM that you want to duplicate.
3. In the High Availability pane, click **Configure**.
4. Enter information for the secondary vGW Security VM. Specify the appropriate IP address information, management network, and data store location.
5. Click **Configure**.



**NOTE:** It is not as important to have backups of vGW Security VMs as it is for the vGW Security Design VM. You can deploy new vGW Security VMs from the templates if necessary.

---

**Related  
Documentation**

- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 11](#)
- [Understanding the vGW Series High Availability Solution on page 3](#)



## PART 2

# Fault Tolerance

- [vGW Series Handling of VMs Enabled for FT on page 19](#)





## CHAPTER 3

# vGW Series Handling of VMs Enabled for FT

- [Understanding vGW Series Fault Tolerance Support on page 19](#)

## Understanding vGW Series Fault Tolerance Support

---

This topic contains the following sections:

- [About vGW Series Fault-Tolerance on page 19](#)
- [vGW Series Fault Tolerance in the vGW Series on page 20](#)
- [Enabling Fault Tolerance for a Virtual Machine on page 20](#)

### About vGW Series Fault-Tolerance

In the virtualized environment, fault-tolerance (FT) ensures continuous support of a virtual machine (VM) in the event of failure of the host on which it resides.

When you enable FT on a VM within VMware vCenter, a copy of the VM, called the secondary VM, is created automatically on another host. The original VM, referred to as the primary VM, and its copy, referred to as the secondary VM (VBM), run in lockstep. If the primary VM's host fails, the secondary VM immediately assumes execution, without loss of connectivity, transactions, or data. For this to occur, the primary VM must be on a host that is part of a cluster of the same kind of hosts with the same configuration. Also, high availability must be enabled on the hosts comprising the cluster.

When you enable the FT feature, the secondary VM is created on a host that is either selected by DRS, if DRS is enabled, or is chosen from any available host in the cluster. The primary VM and the secondary VM have the same name and the same BIOS uuid, but each one has its own vc\_uuid and vi\_id.

The secondary VM has its own .vmx file. Both the primary VM's .vmx file and the secondary VM's .vmx file reside in the same data store directory.

When the primary VM's host fails and the secondary VM takes control, from an external viewpoint it appears as if VMotion had moved the primary VM to the host of the secondary VM and the reverse, that is, as if the secondary VM was moved to the host where the primary VM resided.

## vGW Series Fault Tolerance in the vGW Series

This section explains how the vGW Series handles VMs for which FT is enabled in the vCenter, and how it supports FT overall.

vGW Series handles exposure of FT-enabled VMs to the user in the following ways:

- Secondary VMs are not shown in the VM tree.
- Secondary VMs are not shown in the Machines section of the Settings module.
- In the Settings module Installation section, both the primary VM and the secondary VM are shown in the Secured Network firewall tree. Similar to how the vSphere client marks VMs on an host, the word “secondary” is included after the secondary VM’s name.

For example, a cluster might contain two hosts: host1 and host2. When it was created on host1 in the vCenter, FT was enabled for a VM called my-test-vm. The primary my-test-vm VM remains on host1. A secondary VM called “my-test-vm (secondary)” is created on host2 to support FT. You can view these two VMs in the Settings module Installation section.

- You cannot define a policy for the secondary VM.

vGW is prohibited from reconnecting vNICs and automatically suspending or resuming the VM. To do so would produce undesirable effects. For this reason:

- If a VM has FT configured and it is powered on, you cannot select a VM to secure it using the Setting module Installation section. The check box is grayed out. The tooltip for the VM will show that the VM must be suspended or FT must be disabled before the VM can be secured.
- vGW Series auto-secure feature will not attempt to secure an FT-enabled VM. vGW generates an alert telling you that you must disable FT for that VM or suspend the VM for vGW to secure the VM.

The auto-secure feature monitors for cases in which an FT-enabled VM is disabled and for VMs that are suspended and powered-off. If the VM belongs to an Auto Secure group, then vGW will secure it.

For a VM that has been VMsafe secured for which FT has been enabled, the secondary VM will be created and its VMsafe param0 will be incorrect since it reflects the VC\_uuid of the primary VM rather than its own. However, the vCenter will not try to reconfigure it, since the .vmx of the secondary is read-only and any reconfiguration operation will fail.

## Enabling Fault Tolerance for a Virtual Machine

Before you enable FT for a VM, ensure that High Availability is enabled for the cluster.

To enable FT for a VM in the vCenter:

1. Use the vSphere client to access the vCenter, and locate the host where the VM resides.
2. Right-click the name of the VM.

3. From the displayed menu, select **Fault Tolerance**.
4. Select **Turn On Fault Tolerance**.
5. After reviewing the message noting that DRS automation will be disabled and that the memory reservation of the VM will be changed to the memory size of the VM, accept the changes and click **Yes**.

You c

Verify in the Recent Tasks at the bottom of the page that fault tolerance was turned on for the VM.

**Related Documentation**

- *Understanding vGW Series*



## PART 3

# Index

- [Index on page 25](#)



# Index

## Symbols

#, comments in configuration statements.....	xi
( ), in syntax descriptions.....	xi
< >, in syntax descriptions.....	xi
[ ], in configuration statements.....	xi
{ }, in configuration statements.....	xi
(pipe), in syntax descriptions.....	xi

## B

braces, in configuration statements.....	xi
brackets	
angle, in syntax descriptions.....	xi
square, in configuration statements.....	xi

## C

comments, in configuration statements.....	xi
conventions	
text and syntax.....	x
curly braces, in configuration statements.....	xi
customer support.....	xii
contacting JTAC.....	xii

## D

documentation	
comments on.....	xi

## F

font conventions.....	x
-----------------------	---

## H

high availability.....	3
distributed resource schedule.....	3

## M

manuals	
comments on.....	xi

## P

parentheses, in syntax descriptions.....	xi
--	----

## S

support, technical See technical support	
syntax conventions.....	x

## T

technical support	
contacting JTAC.....	xii

## V

VMware	
high availability.....	3

