



vGW Series Getting Started



Published: 2015-02-19

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
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 Getting Started

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	xi
	Documentation and Release Notes	xi
	Documentation Conventions	xi
	Documentation Feedback	xiii
	Requesting Technical Support	xiv
	Self-Help Online Tools and Resources	xiv
	Opening a Case with JTAC	xiv
Part 1	Basics	
Chapter 1	Getting Started	3
	Understanding vGW Series	3
	Understanding the vGW Series Architecture	4
	Understanding Cloud Computing and vGW Series	6
	Understanding the VMware Infrastructure and vGW Series	7
	Understanding vSphere and the vGW Series	7
	Understanding VMware ESX and ESXi Hosts and the vGW Series	7
	Understanding VMotion and vGW Series	8
	Understanding Hypervisors and vGW Series	8
Chapter 2	SDVM Modules	11
	vGW Security Design VM Modules (VMware)	11
	Understanding the vGW Series Main Module	16
	Dashboard	16
	Status Tab	17
	Events and Alerts Tab	19
	Security Alerts	20
	System Status and Events	21
	Quarantine Tab	22
Chapter 3	Status and Alerts	25
	Understanding vGW Series Status and Alerts	25
	Status	25
	Alerts	25
	E-Mail Alert Settings	26
	SNMP Trap Settings	26
	AutoConfig and Multicast Alerts	26

Part 2	VMware and vGW Series	
Chapter 4	Getting Started	31
	vGW Series Prerequisites and Resource Requirements for the VMware Environment	31
	Overall Resource and Access Requirements	31
	Virtual Appliance System Requirements	32
	vGW Series VMware vSwitch Requirements	33
	VMware Port Group Requirements	34
	Virtualized NIC Requirements	35
	Preparing to Integrate the vGW Series with the VMware Environment	35
	Understanding vGW Series Environment Time Synchronization	36
	vGW Series VMsafe Firewall + Monitoring and VMsafe Monitoring Modes	36
Chapter 5	OVA and vGW Series Deployment	39
	Understanding the Open Virtualization Format OVA Template Method	39
	Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure	40
	Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware	49
	Using the OVA Single File Method to Integrate the vGW Security VM with VMware	51
Part 3	vGW Series Setup	
Chapter 6	SDVM Set Up Process	55
	Setting Up vGW Series	55
Part 4	SDVM and SVM	
Chapter 7	Basics	63
	Understanding the vGW Security Design VM	63
	Understanding the vGW Security VM	64
	Understanding the vGW Series Kernel Module	64
Chapter 8	SDVM Navigation and VM Tree	67
	Understanding vGW Security Design VM Navigation	67
	Understanding the vGW Security Design VM Taskbar	69
	About the vGW Security Design VM Tree	70
	VM Tree Overview	71
	Locating VMs in a Complex VM Tree	72
Part 5	vGW Series IPv6 Support	
Chapter 9	IPv6 Addressing	77
	Understanding IPv6 Addressing	77
	IPv6 and the Cloud	77
	IPv6 and IPv4	78
	IPv6 Address Space, Addressing, and Address Types	78
	The IPv6 Basic Packet Header	78
	The IPv6 Packet Header Extensions	80

	The IPv6 Address Format	81
	Address Assignment and IPv6	81
	Understanding vGW Series IPv4 and IPv6 Dual Stack Support	82
	Dual Stack Background	82
Chapter 10	vGW Series IPv6 Implementation	85
	Overview of IPv6 Implementation in the vGW Security Design VM Modules	85
	Main Module	85
	Network Module	86
	Firewall Module	86
	Firewall Logs	86
	Policies	86
	ICMPv6	86
	IDS Module	87
	AntiVirus Module	87
	Introspection Module	87
	Compliance Module	87
	Reports Module	87
	Settings Module	87
	Understanding vGW Series IPv6 Support	87
	vGW Security Design VM and vGW CLI Support for IPv6 Addresses	88
	Entering IPv6 Addresses	89
	vGW Series IPv6 Address Representation	89
	IPv4-Mapped IPv6 Addresses	90
	vGW Security Design VM Filter Boxes	90
	Searching the VM Tree for VMs and Hypervisors with IPv6 Addresses	90
	vGW Security Design VM and IPv6 and IPv4 Addressing	90
	vGW Security VM IP Addressing Support	90
	IPv6 Support in Homogeneous and Heterogeneous vGW Series	
	Environments	91
	IPv6 Traffic Handling in Homogenous Environments (All vGW Series	
	Components at Version 5.5 or Later)	91
	IPv6 Traffic Handling in Heterogeneous Environments (with a Mix of vGW	
	Series Component Versions)	91
Chapter 11	IPv6 and vGW Series Requirements	95
	Understanding Requirements for Communication Between vGW Series IPv6	
	Components and Juniper Networks IPv4 Servers	95
	Understanding vGW Series IPv6 Accessibility Requirements for Customer	
	Servers	96
Part 6	Index	
	Index	99

List of Figures

Part 1	Basics	
Chapter 1	Getting Started	3
	Figure 1: vGW Series Architecture	5
Chapter 2	SDVM Modules	11
	Figure 2: Main Module	12
	Figure 3: Network Module	12
	Figure 4: Firewall Module	13
	Figure 5: IDS Module	13
	Figure 6: AntiVirus Module	14
	Figure 7: Introspection Module	14
	Figure 8: Compliance Module	15
	Figure 9: Reports Module	15
	Figure 10: Settings Module	16
	Figure 11: Dashboard Tab	17
	Figure 12: Status Tab	18
	Figure 13: Taskbar Showing the Health Status Icon	19
	Figure 14: Main Module Events and Alerts Page	20
	Figure 15: Consolidated Logs for Events and Alerts	20
	Figure 16: Quarantine Tab	22
Part 2	VMware and vGW Series	
Chapter 5	OVA and vGW Series Deployment	39
	Figure 17: OVA Template Details Page	42
	Figure 18: OVA File Deployment License Agreement	42
	Figure 19: Naming the vApp	43
	Figure 20: Specifying the Host and Cluster	44
	Figure 21: Selecting the Storage	44
	Figure 22: Mapping the vGW Management Networks	45
	Figure 23: Specifying the Database Disk Size	46
	Figure 24: Verifying That the Configuration Is Correct	46
	Figure 25: Displaying the vGW Appliance Components	47
	Figure 26: vGW Security Design VM Summary Tab in vCenter	49
Part 3	vGW Series Setup	
Chapter 6	SDVM Set Up Process	55
	Figure 27: vGW Series Security Design VM Login Screen	56
	Figure 28: vGW Series Licensing	58
	Figure 29: Authenticating to the vGW Installation Wizard	58

Part 4	SDVM and SVM	
Chapter 7	Basics	63
	Figure 30: Main Module Displayed at Login	64
Chapter 8	SDVM Navigation and VM Tree	67
	Figure 31: vGW Security Design VM Taskbar	67
	Figure 32: VM Tree	68
	Figure 33: vGW Security Design VM Taskbar	69
	Figure 34: VM Tree with Selected VMs	71
	Figure 35: Searching All VMs in the VM Tree Using the Advanced Editor	73
	Figure 36: Searching for Specific VMs in the VM Tree Using the Advanced Editor	74
Part 5	vGW Series IPv6 Support	
Chapter 10	vGW Series IPv6 Implementation	85
	Figure 37: Console Showing IPv6 and IPv4 vGW Security VM Addresses	88

List of Tables

	About the Documentation	xi
	Table 1: Notice Icons	xii
	Table 2: Text and Syntax Conventions	xii
Part 1	Basics	
Chapter 2	SDVM Modules	11
	Table 3: vGW Series Status Icons	18
Part 4	SDVM and SVM	
Chapter 8	SDVM Navigation and VM Tree	67
	Table 4: Taskbar Icons	69
	Table 5: Virtual Machine State Icons	72
Part 5	vGW Series IPv6 Support	
Chapter 9	IPv6 Addressing	77
	Table 6: IPv6 Basic Packet Header Fields	78
	Table 7: IPv6 Extension Headers	80

About the Documentation

- Documentation and Release Notes on page xi
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiv

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 xii 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 xii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<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@# set system domain-name <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">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(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.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
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">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

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. 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

Basics

- [Getting Started on page 3](#)
- [SDVM Modules on page 11](#)
- [Status and Alerts on page 25](#)

CHAPTER 1

Getting Started

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Architecture on page 4](#)
- [Understanding Cloud Computing and vGW Series on page 6](#)
- [Understanding the VMware Infrastructure and vGW Series on page 7](#)
- [Understanding Hypervisors and vGW Series on page 8](#)

Understanding vGW Series

vGW Series delivers complete virtualization security for multitenant public and private clouds, and clouds that are a hybrid of the two. vGW Series comprises the following three main components:

- The vGW Security Design VM that provides a central management server. It consists of a set of modules that you use to configure the vGW Series features for your virtualized environment. It provides charts, tables, and graphs that allow you to view information that vGW Series produces about your environment and use in determining how to adjust your security policy.

You use it to install and manage the vGW Security VMs that you deploy to secure hosts in your virtualized environment.

- The vGW Security VM that is installed on each host to be secured. The vGW Security VM acts as a conduit to the vGW kernel module that it inserts into the hypervisor of the host that vGW Series protects. The vGW Security VM maintains policy and logging information. A vGW Security VM remains attached to the ESX/ESXi host that it is installed on.

The vGW Security Design VM pushes the appropriate security policy to the vGW Security VM which, in turn, inserts it into the vGW kernel module.

- The vGW kernel module

Virtualized network traffic is secured and analyzed against the security policy for all VMs on the ESX/ESXi host in the vGW kernel module installed on the host. All connections are processed and firewall security is enforced in the vGW Series kernel module.

**Related
Documentation**

- [Understanding Cloud Computing and vGW Series on page 6](#)
- [Understanding Hypervisors and vGW Series on page 8](#)
- [Understanding the vGW Security Design VM on page 63.](#)
- [Understanding the vGW Security VM on page 64](#)

Understanding the vGW Series Architecture

vGW Series is a fault-tolerant service provider and enterprise grade security solution that is purpose-built for the virtualized environment. Not only does it secure virtual machines (VMs), but it also protects the hypervisor. When it is deployed into the VMware environment and the vGW Security VM is installed on a VMware ESX/ESXi host, the vGW kernel module (vGW engine) is loaded into the host's hypervisor between the virtual network installation card (vNIC) and the virtual switch (vSwitch). The VMware VMsafe module gives vGW Series full protocol inspection of every VM.

vGW Series does not depend on the virtual switching layers for its oversight of VMs. Consequently, whichever vSwitch is used has no bearing on vGW Series. It is compatible with them all.



NOTE: VMware lets you create abstracted network devices called virtual switches (vSwitches). A vSwitch routes traffic internally between virtual machines and it links to external networks. A vSwitch can be connected to physical switches.

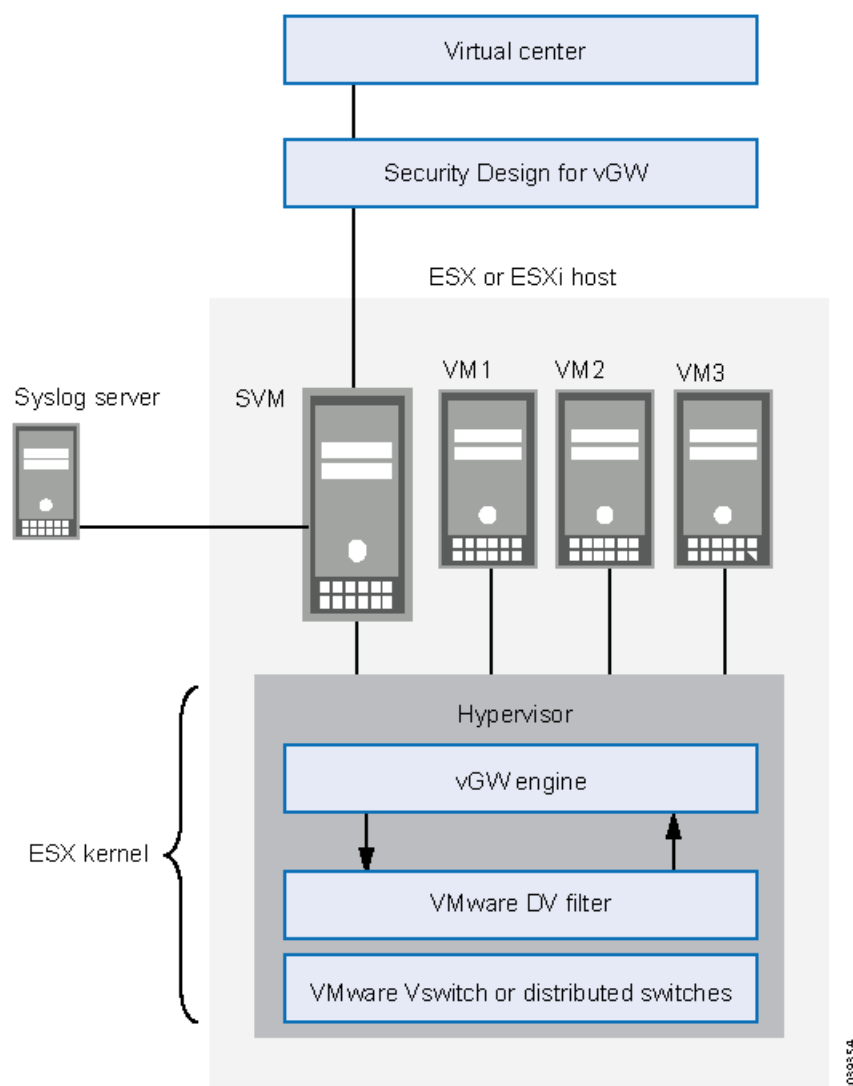
In the VMware virtualized environment, port groups are used to aggregate multiple ports under a common configuration. They serve as an anchor point for virtual machines that connect to labeled networks.

The vGW Security Design VM makes configuration changes in the VMware vCenter automatically. This lowers administrative complexity and reduces the possibility of configuration errors. [Figure 1 on page 5](#) shows the vGW Series integration with VMware ESX/ESXi hosts and vCenter. [Figure 1 on page 5](#) also shows that:

- The vGW Security Design VM is integrated and communicating with the VMware vCenter. It is also communicating with the vGW Security VM installed on the ESX/ESXi host.
- The vGW Security VM, which is installed on the host, has inserted the vGW engine (the vGW kernel module) into the ESX/ESXi host's hypervisor.

From within the ESX/ESXi host's hypervisor, the vGW Series engine is aware of all network connections between VMs on the host, coming into and going out to other hosts in the virtualized environment, and transiting the physical switch.

Figure 1: vGW Series Architecture



Related Documentation

- [Configuring the vGW Series to Send Syslog and Netflow Data to Juniper Networks STRM](#)
- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability](#)
- [Installing a Secondary vGW Security VM for High Availability](#)
- [Integrating the vGW Series with VMware Using the Settings Module](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 35](#)
- [Understanding Hypervisors and vGW Series on page 8](#)
- [Understanding the vGW Security VM on page 64](#)
- [Understanding the vGW Series High Availability Solution](#)
- [Understanding the vGW Series Hypervisor and Extended VM Security](#)

- *Understanding vGW Series Fault Tolerance Support*
- *Viewing the vGW Series Logs*

Understanding Cloud Computing and vGW Series

A cloud is an Internet-based environment of virtualized computing resources including servers, software, and applications that can be accessed by individuals or businesses with Internet connectivity. Customers, referred to as tenants, can access the resources that they need to run their businesses.

Clouds possess the following advantages. They:

- Allow customers to share the same infrastructure to gain price and performance advantages.
- Provide customers with a pay-as-you-go lease-style investment versus buying all of the required hardware and software up front themselves.
- Allow businesses to scale easily and tier more services and functionality on an as-needed basis.

There are two kinds of clouds and a third one that combines the other two:

- Public clouds

They are based on a standard cloud computing model. In this structure, a service provider (SP) hosting the cloud makes resources such as applications, computing capacity, storage, and server-based infrastructure available to the public.

The SP hosting the cloud owns and operates the infrastructure and offers access through the Internet.

- Private clouds

They are proprietary network or data center that use cloud computing technologies such as virtualization.

The infrastructure is operated solely for a single organization whether managed internally or externally. The company still must buy, build, and manage the infrastructure. It does not benefit from the economic gains offered by public cloud computing.

- Hybrid clouds

They are composed of two or more clouds that remain unique but are bound together providing benefits of multiple deployment models. They are maintained by both internal and external providers.

Hybrid clouds require both on-premises resources and off-site, remote server-based cloud infrastructure.

Cloud computing allows for dynamic and elastic generation of virtual infrastructure and virtual machines (VMs) with their own operating systems running over that infrastructure.

Whether for public, private, or hybrid clouds, virtualized data centers must offer secure, discrete, VM environments to their customers and their organizations.

Physical network security is designed for and limited to physical hardware and its software. It does not have the visibility into traffic transmission and communication between VMs that is required to secure the environment. vGW Series secures the virtual network in ways that physical security mechanisms protecting physical networks cannot because it is purpose-built for virtualized environments. vGW Series provides support for IPv6 in addition to IPv4 to enable organizations that have adopted IPv6 to benefit from vGW Series cloud security.

Related Documentation

Understanding the VMware Infrastructure and vGW Series

The Juniper Networks vGW Series runs as integrated software on VMware vSphere servers.

This topic includes the following sections:

- [Understanding vSphere and the vGW Series on page 7](#)
- [Understanding VMware ESX and ESXi Hosts and the vGW Series on page 7](#)
- [Understanding VMotion and vGW Series on page 8](#)

Understanding vSphere and the vGW Series

VMware vSphere is a cloud operating system that can manage large pools of virtualized computing infrastructure, including software and hardware. vGW Series components integrate with the VMware vSphere infrastructure to provide security for ESX/ESXi hosts in the virtualized environment. Because the vGW Series is purpose-built to support virtualization, it synchronizes automatically with the VMware vCenter. It uses VMware's *VMsafe* interfaces to provide breakthrough levels of security and performance.



NOTE: Beginning with vGW Series 5.0r2, vGW Series provides support for vSphere 5.0.

Understanding VMware ESX and ESXi Hosts and the vGW Series

VMware ESX and ESXi hosts provide the foundation for building and managing a virtualized IT environment. These hypervisor-based hosts contain abstract processors, memory, storage, and networking resources that are shared among multiple virtual machines (VMs) that run unmodified, diverse operating systems and applications.

vGW Series manages and secures the VMs that run on ESX/ESXi hosts.

The number of IP addresses or VMs that vGW Series can protect is not determined. In any case, a single vGW Security Design VM management center can handle hundreds of hosts and their associated vGW Security VMs, and each vGW Security VM can load thousands of policy rules. However, a vGW Security VM loads only the policy rules that

are relevant for the VMs which exist on the host where it resides. You can easily extend the reach of protection for your virtualized environment, if it is exceedingly large, by using the vGW Series Split Center and Multi-Center features, which allow you to scale to accommodate any size requirements.

Understanding VMotion and vGW Series

VMware provides a feature called *VMotion* that allows for transition of active, or live, VMs from one physical server to another. VMs can be moved from one server to another to perform maintenance operations on a host. Also, they can be moved automatically when VMotion is triggered through VMware's Dynamic Resource Scheduler (DRS), which is used to evenly distribute system resource usage across physical servers.

Because VMs can be migrated between servers, their security levels can be compromised and lowered from that of the original server to that of the new one. A VM could be migrated to an unsecured zone or one with a lower trust level.

Unlike traditional firewalls, the vGW Series firewall supports live migration by maintaining open connections and security throughout the event. vGW Series ensures that appropriate security for a VM remains intact throughout migration.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the Open Virtualization Format OVA Template Method on page 39](#)
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)
- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 40](#)
- [Integrating the vGW Series with VMware Using the Settings Module](#)
- [Understanding the vGW Series High Availability Solution](#)
- [Understanding vGW Series Fault Tolerance Support](#)
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)

Understanding Hypervisors and vGW Series

vGW Series is a high performance, hypervisor-based virtualization security solution. Various layers of abstraction combine to create virtualized environments. Virtualized hardware supports multi-tenancy in which guest virtual machines (VMs) running discrete operating system images share the system resources. Each guest VM runs its own user-space applications. If a physical machine does not directly support virtualization, a software layer called a hypervisor is used to manage the relationship between the guest VMs that run on it and compete for its resources.

Some forms of virtualization provide support for multiple instances of the same kind of guest operating system. A full-virtualization hypervisor allows multiple instances of a *variety* of guest operating systems to run concurrently. It presents a virtual operating

platform to the guest operating systems, and it manages their execution. Similar to the control program responsibilities of the supervisor that is intrinsic to some operating systems, a full-virtualization hypervisor arbitrates access to resources between guest VMs that reside on the host and it manages those guest VM operating systems.

vGW Series secures full virtualized environments by inserting a module into the hypervisor of the host. Because the vGW Series kernel module resides in the hypervisor, vGW Series has wide visibility into the virtualized environment, and it can process security requirements faster.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Understanding Cloud Computing and vGW Series on page 6](#)
- [Understanding vGW Series IPv6 Support on page 87](#)

CHAPTER 2

SDVM Modules

- [vGW Security Design VM Modules \(VMware\) on page 11](#)
- [Understanding the vGW Series Main Module on page 16](#)

vGW Security Design VM Modules (VMware)

The vGW Security Design VM is composed of the following modules that implement vGW Series features:

- Main
- Network
- Firewall
- IDS
- AntiVirus
- Introspection
- Compliance
- Reports
- Settings

The following figures show the modules' primary pages. A link is provided to the section that covers the module. The highlighted button on the taskbar at the top of the page indicates the active feature.

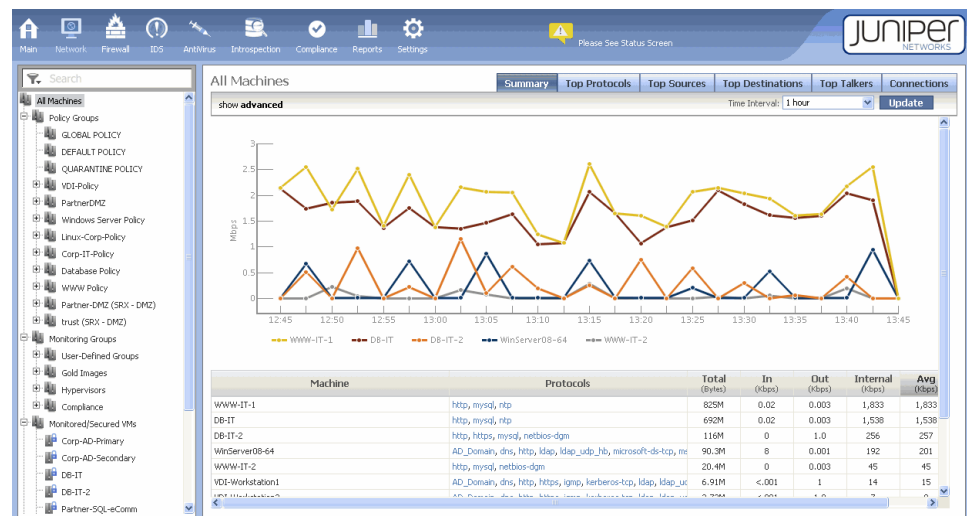
- Main. See “Understanding the vGW Series Main Module” on page 16 and Figure 2 on page 12.

Figure 2: Main Module



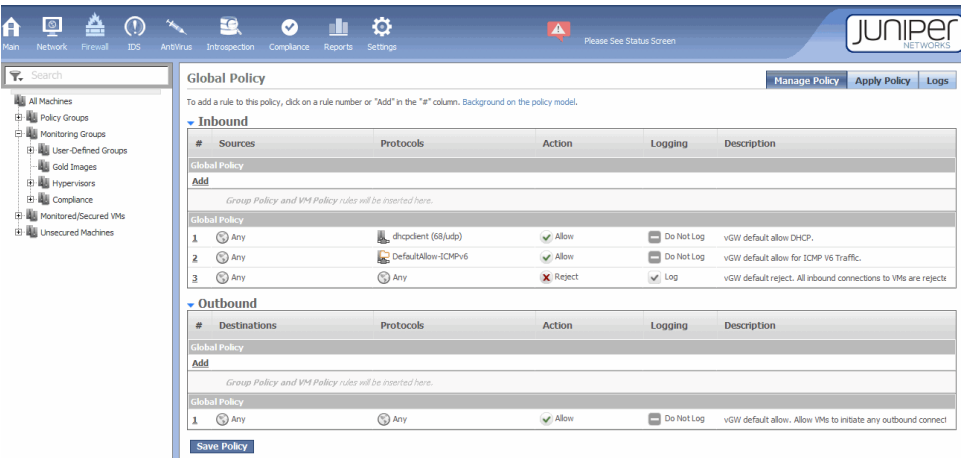
- Network. See *Understanding the vGW Series Network Module* and Figure 3 on page 12.

Figure 3: Network Module



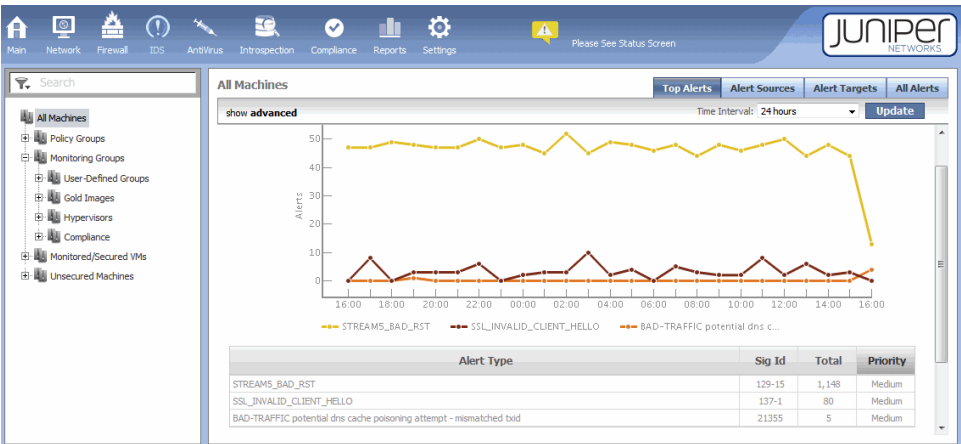
- Firewall. See *Understanding the vGW Series Firewall Module*. See Figure 4 on page 13.

Figure 4: Firewall Module



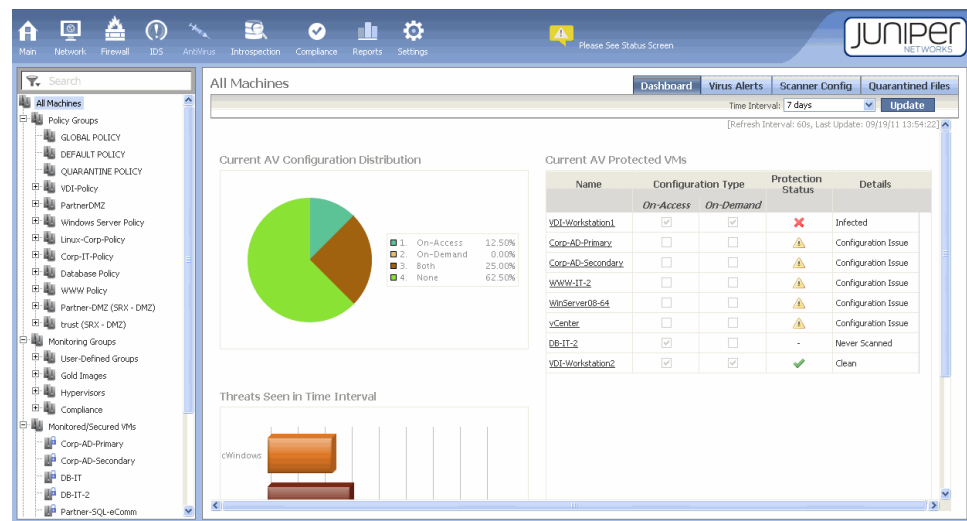
- IDS. See *Understanding the vGW Series IDS Module* and Figure 5 on page 13.

Figure 5: IDS Module



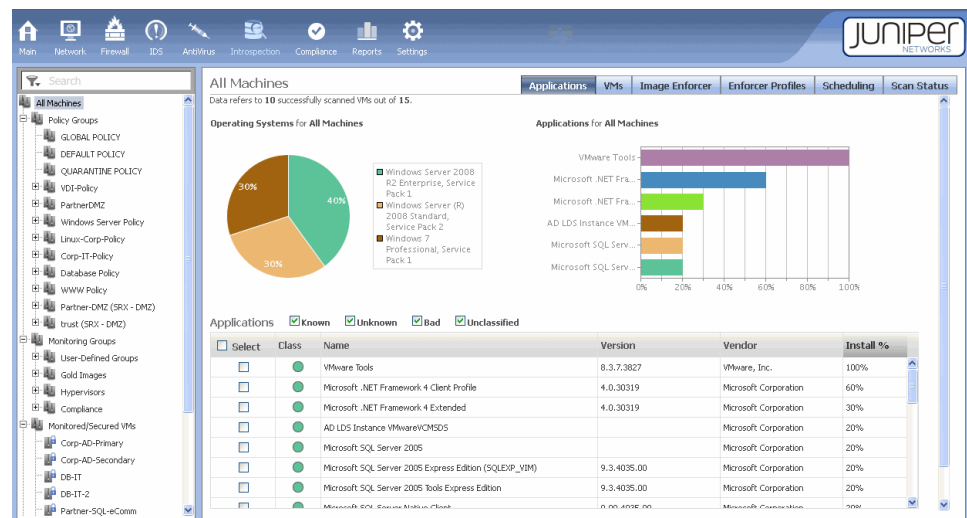
- AntiVirus. See *Understanding vGW Series AntiVirus* and Figure 6 on page 14.

Figure 6: AntiVirus Module



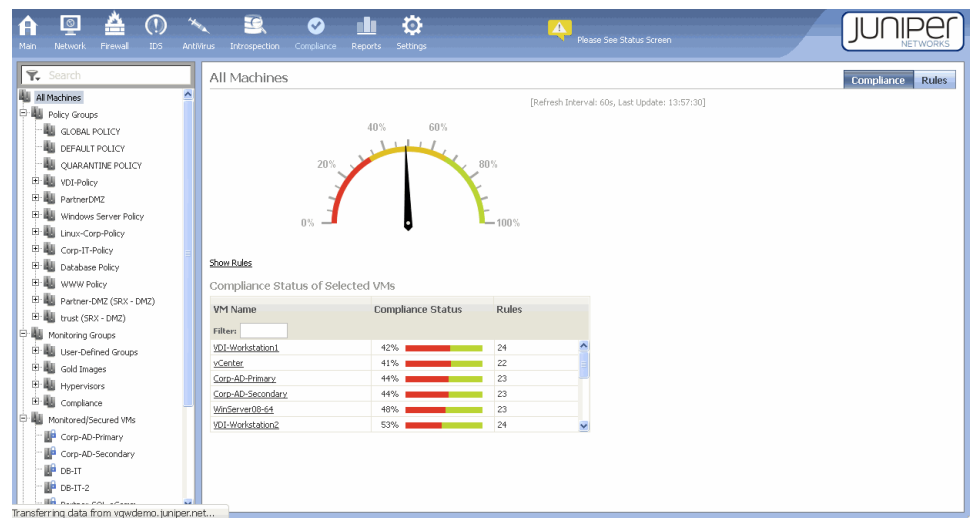
- Introspection. See *Understanding the vGW Series Introspection Module* and Figure 7 on page 14.

Figure 7: Introspection Module



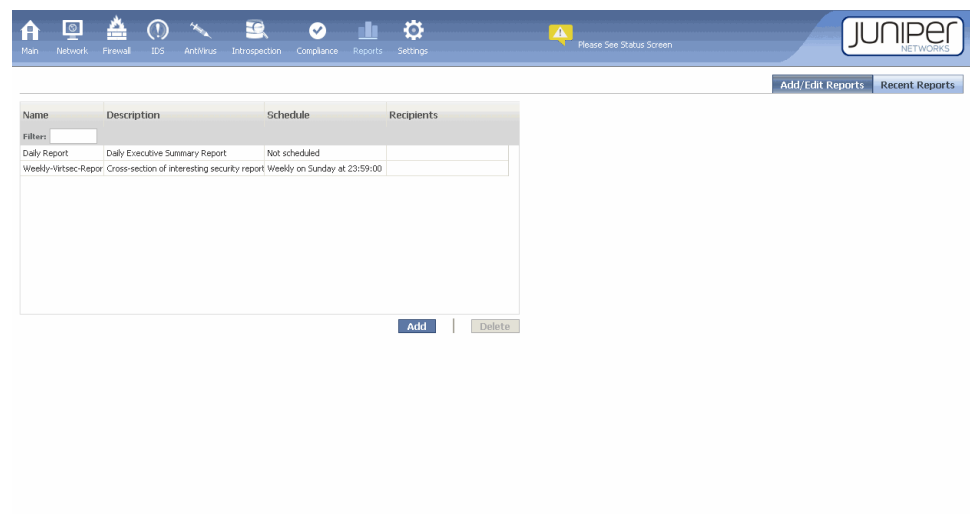
- Compliance. See *Understanding the vGW Series Compliance Module* and Figure 8 on page 15.

Figure 8: Compliance Module



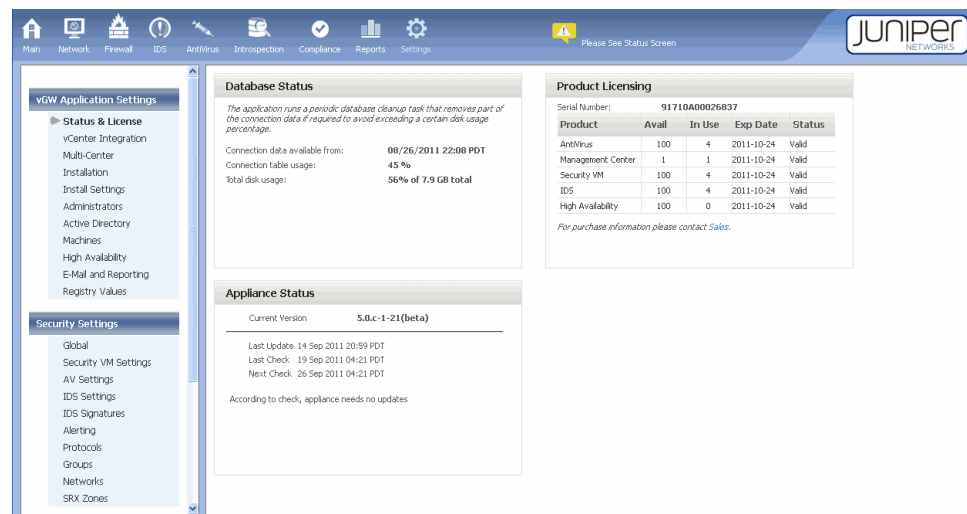
- Reports. See *Understanding the vGW Series Reports Module* and [Figure 9 on page 15](#).

Figure 9: Reports Module



- Settings. See *Understanding the vGW Series Settings Module* and [Figure 10 on page 16](#).

Figure 10: Settings Module



Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding Cloud Computing and vGW Series on page 6](#)
- [Understanding Hypervisors and vGW Series on page 8](#)
- [Understanding the vGW Security VM on page 64](#)

Understanding the vGW Series Main Module

The Main module of the vGW Security Design VM displays information gathered from many of the vGW Security Design VM components. When vGW Series detects new events and alerts, data and graphs in the Main module's panes are automatically refreshed.

The Main module contains the following tabs.

- [Dashboard on page 16](#)
- [Status Tab on page 17](#)
- [Events and Alerts Tab on page 19](#)
- [Quarantine Tab on page 22](#)

Dashboard

In both graphical and table format, the Dashboard allows you to view the behavior of your environment at a glance. You can view the activity of all virtual machines (VMs). You can select an individual VM or a group of VMs in the VM tree to focus on. The Dashboard displays information for both IPv4 and IPv6 traffic.

See [Figure 11 on page 17](#).

Figure 11: Dashboard Tab



The Dashboard includes the following panes:

vGW Status—Provides an overview of the current state of your infrastructure. It shows the state of vGW connectivity to the VMware vCenter. It also shows the number of vGW Security VMs deployed to secure ESX/ESXi hosts, and the overall state of your deployment's VMs, that is, whether they are secured by vGW Series or not.

Compliance Status for All Machines—Shows the overall posture of all VMs in your organization that might be violating compliance rules. The more VMs that violate rules (high weighting), the further the needle moves to the red.

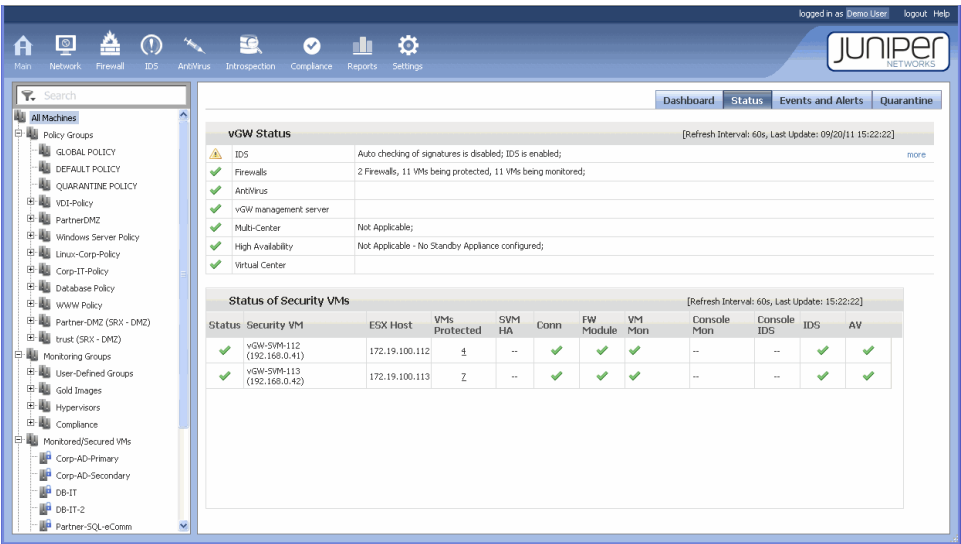
Top Talkers for All Machines —Displays network activity for the last hour.

IDS Alerts for All Machines—If IDS is enabled, the overall IDS alerts information is displayed.

Status Tab

The Status tab displays a summary of vGW status for each module, and it displays status on individual vGW Security VMs. The page is refreshed every 60 seconds. See [Figure 12 on page 18](#).

Figure 12: Status Tab



NOTE: For vGW Security VMs for which standby or secondary vGW Security VM instances are configured, vGW Series counts only the primary vGW Security VM and reflects that count in vGW Status table Firewalls number.

For disconnected vGW Security VMs, Firewalls shows separate counts for primary, standby, and secondary vGW Security VMs. For example, it might show “1 disconnected, 1 Standby disconnected, 1 Secondary disconnected”.

The Status page includes these panes:


vGW Status—For the vGW Series components, the pane indicates the current state using the status icons shown in [Table 3 on page 18](#).

Table 3: vGW Series Status Icons

Icon	Indicates
	vGW Series component is working properly.
	One or more issues exist with the component. For example, maintenance settings might be incompatible or disabled, or you might need to update its firewall.
	Significant issues exist for the component. For example, a module did not load correctly.

In addition to these icons, an overall health status icon appears when individual components require your attention. [Figure 13 on page 19](#) shows the taskbar with the health status icon at the far right. The icon is either red or yellow, depending on the underlying state of the components being monitored.

Figure 13: Taskbar Showing the Health Status Icon



Status of Security VMs—This pane reports status on individual vGW Security VMs.

This pane shows the following information:

- vGW Security VM name.
- Host that the vGW Security VM protects.
- Number of VMs that it protects.

For vGW Security VMs configured with secondary or standby instances, vGW Series counts VMs protected by the primary vGW Security VM. That is, it does not count the same VM again in relation to the secondary or the standby vGW Security VM instance.

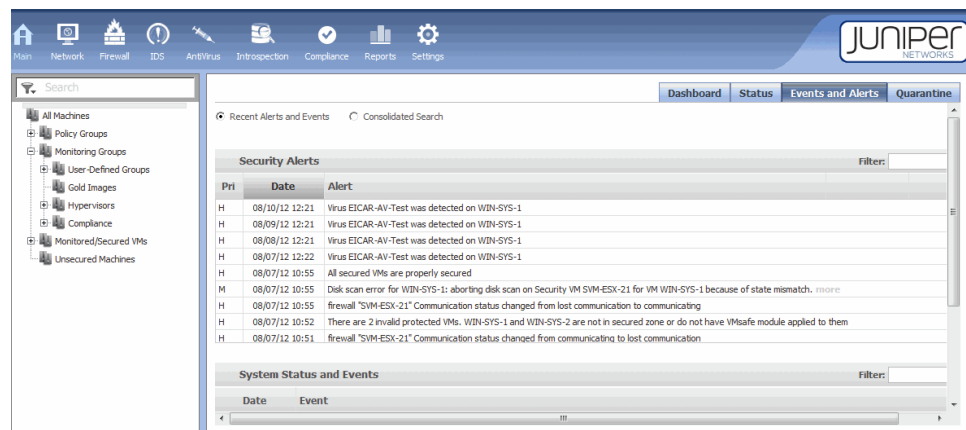
- If vGW Series HA is enabled.
- If the vGW Security VM is connected to the vGW Security Design VM.
- If the firewall module is enabled.
- If VM monitoring is used.
- IP address of the vGW Security Design VM management center.
- If IDS is used, the IP address of the IDS console.
- If IDS is enabled, IDS data appears. Otherwise, the chart is blank.
- If AntiVirus is enabled.

Click the Status icon for a vGW Security VM to display detailed information about it. When you click the icon, the vGW Security Design VM automatically positions you in the Security VM Settings section of the Settings module that pertains to the selected vGW Security VM. You can use the tabs on that page to change configuration settings for the vGW Security VM. See *Understanding the vGW Security VM Settings*.

Events and Alerts Tab

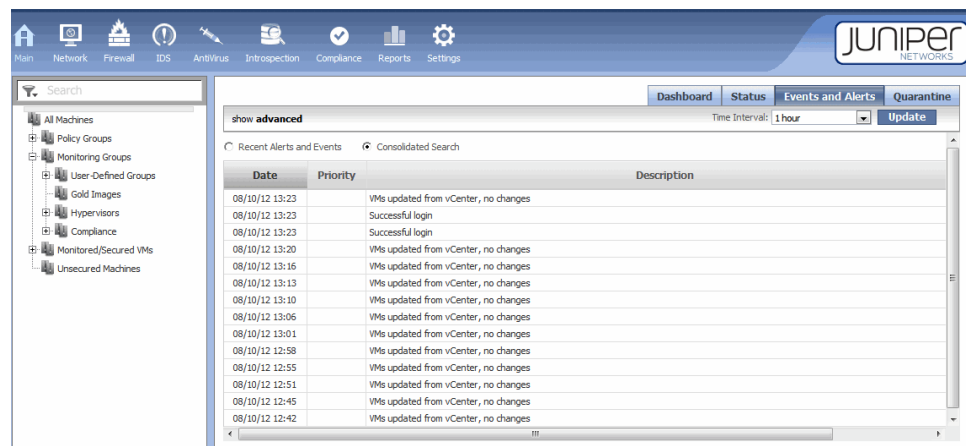
The Events and Alerts page allows you to view Security Alerts and System Status and Events messages individually, in separate panes of the page. You can use an individual filter to search each set separately. See [Figure 14 on page 20](#).

Figure 14: Main Module Events and Alerts Page



Alternatively, you can search through the combined logs for a specific time period using the Consolidated Search button. For example, you might want to look at historical data. Rather than searching through each set, you can specify a time and see all the logs for that period. See [Figure 15 on page 20](#).

Figure 15: Consolidated Logs for Events and Alerts



- [Security Alerts on page 20](#)
- [System Status and Events on page 21](#)

Security Alerts

The Security Alerts pane lists all vGW Series alerts that have occurred in your protected virtualized environment, except for IDS alerts and AntiVirus alerts which are reported in their own modules. The reported alerts are primarily vGW Series system-related events, such as reports on occurrences of vGW Series version updates or alerts when component failures occur.

Alerts are classified as high (H), medium (M), or low (L), depending on their severity. Click the **Priority** or **Date** column to sort the list differently. You can use the filter to sort the data by IPv6 or IPv4 address. The pane will show the alert or event for only the VM with the IP address that you enter.

System Status and Events

Many companies require a complete audit trail of administrative and policy operations to meet compliance standards and their security best practices. A detailed audit trail is an important part of a security infrastructure that security administrators rely on.

vGW Series collects information on events and posts it to the System Status and Events pane when administrative and policy operations occur. It posts the following event alerts:

- An administrator logs in or logs out, and when failed login attempts occur.
- License changes are made.
- An administrator changes vGW Security Design VM settings, including the following:
 - Changes to general system settings such as log connections, system reboots, and active directory.
 - Manual VM updates.
 - Modifications to vGW Series objects, including networks, machines, groups, protocols, an administrator settings.
 - Updates to the vGW Security Design VM.
 - Updates to the vGW Security VM.
 - Configuration changes to firewall.
 - Configuration changes to Syslog, Netflow, external inspection devices, and infrastructure reinforcement.
- Automatically secured VM configuration changes occur.
- IDS signatures are modified and new signatures are added.
- Introspection scans are started on **Scan Now** requests, scheduled events occur, and scheduled scan configurations are modified.
- Compliance Rule modifications are made.
- Reports are created or Reports configuration settings are modified.
- The Image Enforcer is configured, its configuration settings are changed, and Image Enforcer scans occur.
- AntiVirus is configured, changes are made to its configuration, and AntiVirus scans occur.
- SRX Series integration changes take place.
- Multi-Center and Split-Center settings are configured or changed.
- Backup and Restore is configured and when configuration changes are made.
- Registry values are changed.

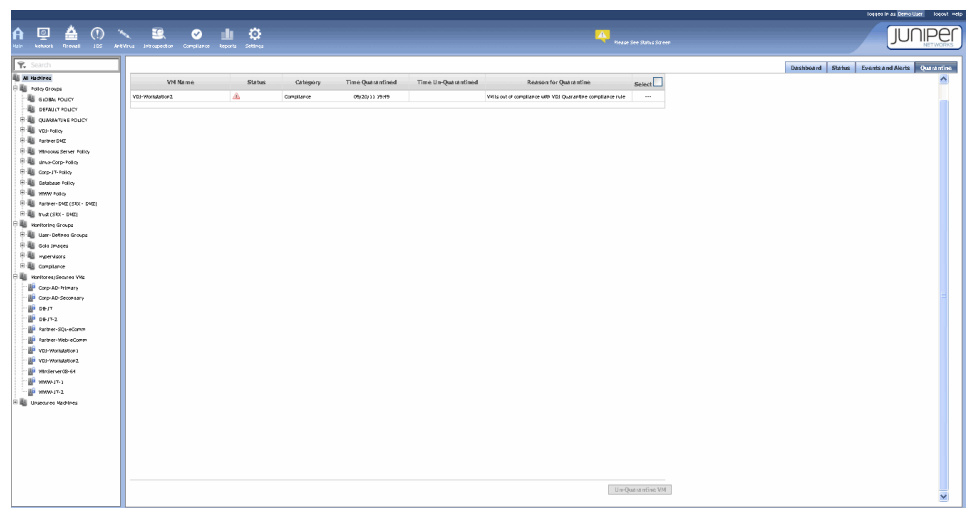
Events are listed chronologically. The events that occurred most recently are listed at the top of the table. To view additional events, you can access the vGW Security Design VM database.

You can configure the Alerting pane in the Settings module to allow alerts to be sent also to administrators through e-mail. See *vGW Series Event and Alert Messages Guide Reference*.

Quarantine Tab

The Main module Quarantine tab displays information about VMs that have been quarantined as a result of AntiVirus, Compliance, or Image Enforcer scans. Using it, you can view the time that the VM was quarantined, when it was removed from quarantine, and the reason that it was quarantined. You can also remove a VM from quarantine from this page. See [Figure 16 on page 22](#).

Figure 16: Quarantine Tab



To display information about quarantined VMs for one or more features, select the check box beside the feature. You can view information about VMs quarantined as a result of only one type of scan or you can view all information for any of them in combination. For any of these selections you can display:

- Information about currently quarantined VMs.
- Historical information about previously quarantined VMs.

The Quarantine page shows the following information for each VM:

- Status
- Category
- Time quarantined
- Time un-quarantined.
- Reason why the VM was quarantined.

To remove a VM from quarantine, check the select box for it and click **Un-Quarantine VM**.



NOTE: You can use the AntiVirus module to quarantine files infected by a virus or other malware. See *Understanding vGW Series AntiVirus*.

For details on the relationship between the Main module Quarantine tab, the Quarantine Policy group, and AntiVirus, Compliance, and Image Enforcer scans, see *Understanding Quarantined VMs and How to Manage Them*.

**Related
Documentation**

- [Understanding the vGW Security Design VM on page 63](#)
- [Understanding the vGW Security Design VM Taskbar on page 69](#)
- [About the vGW Security Design VM Tree on page 70](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 3

Status and Alerts

- [Understanding vGW Series Status and Alerts on page 25](#)

Understanding vGW Series Status and Alerts

vGW Series can display several status icons within the user interface and several mechanisms for sending alerts, so that you know exactly what is happening on the virtual network.

- [Status on page 25](#)
- [Alerts on page 25](#)
- [E-Mail Alert Settings on page 26](#)
- [SNMP Trap Settings on page 26](#)
- [AutoConfig and Multicast Alerts on page 26](#)

Status

vGW Series interface displays a yellow or red status icon to indicate an event or configuration issue that merits attention.

Click the status icon to display the Status tab in the Main module's page.

The sections of the product that have triggered a status change are displayed with most important status changes at the top shown in red. For details on the status issues, click the more link next to the status summary line.

Alerts

vGW Series can send alerts when the log field in a rule in a security policy is set to Alert or Custom E-Mail Alert Tag and a connection matching this rule is seen on the network.

In addition to alerts generated by security rules, vGW Series monitors High, Medium and Low Security events, displayed on the Main module's Events and Alerts tab, and it reports those Alerts out through the settings here (that is, through E-Mail, SNMP trap, or both).

In both cases, alerts use the settings found in Settings -> Security Settings -> Alerting.

You can choose to send an e-mail alert and an SNMP trap, only e-mail alerts, or only SNMP traps.

E-Mail Alert Settings

Enable e-mail alerts by providing the mail relay server IP address as well as the source and destination e-mail addresses. The aggregation time is the gap between successive notifications.

You are not required to configure multiple e-mail recipients. However, four custom e-mail alert tags can be created that point to different e-mail aliases or individual e-mail accounts (or a combination of the two). These custom tags can then be specified in the security policy editor.

If you want to send both an e-mail alert and an SNMP trap on a single rule, you can do so by using the standard alert icon. However, only the e-mail addresses listed in the Recipients Addresses are used. In other words, custom tags cannot be used when sending e-mail and SNMP alerts.

SNMP Trap Settings

Simple Network Management Protocol (SNMP) is an IP protocol used mostly to monitor network-attached devices for conditions that warrant administrative attention. SNMP exposes management data in the form of variables on the managed systems, which describe the system configuration. These variables can then be queried (and sometimes set) by managing applications. In typical SNMP uses, one or more administrative computers, called managers, have the task of monitoring or managing a group of hosts or devices on a computer network. Each managed system executes, at all times, a software component called an agent which reports information via SNMP to the manager. SNMP operates in the Application Layer of the Internet Protocol Suite (Layer 7 of the OSI model). The SNMP agent receives requests on UDP port 161. The manager may send requests from any available source port to port 161 in the agent. The agent response will be sent back to the source port on the manager. The manager receives notifications (Traps and InformRequests) on port 162

SNMP traps can be set through SNMPv1 or SNMPv2. You must enter the SNMP server address and community string. You can again set the aggregation time (the delay between successive events), if desired.

AutoConfig and Multicast Alerts

By default the vGW Series is configured to alert when autoconfig addresses are discovered (Settings -> Security Settings -> Alerting). No alert is automatically sent when Multicast is seen (though this can be enabled).

- Autoconfig addresses: When a machine does not have an IP address configured or cannot acquire a DHCP lease, it defaults to an autoconfig address in the 169.254.** range. This setting often represents a configuration problem or an issue with the DHCP service.
- Multicast: Many hosts use multicast packets to advertise their presence on the network as well as broadcast information regarding which services they offer and configuration data. This information is often not needed, so it can be undesirable for servers to provide

it. In addition, there are security issues related to advertising the services a machine has available.

Related Documentation • [Understanding vGW Series on page 3](#)

PART 2

VMware and vGW Series

- [Getting Started on page 31](#)
- [OVA and vGW Series Deployment on page 39](#)

CHAPTER 4

Getting Started

- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 35](#)
- [Understanding vGW Series Environment Time Synchronization on page 36](#)
- [vGW Series VMsafe Firewall + Monitoring and VMsafe Monitoring Modes on page 36](#)

vGW Series Prerequisites and Resource Requirements for the VMware Environment

This topic covers how to prepare to install the vGW Series product for integration and deployment in the VMware vSphere environment. It covers prerequisites and identifies the resources required to import the vGW Series into the VMware environment, install the product, and run it.

This topic includes the following sections:

- [Overall Resource and Access Requirements on page 31](#)
- [Virtual Appliance System Requirements on page 32](#)
- [vGW Series VMware vSwitch Requirements on page 33](#)
- [VMware Port Group Requirements on page 34](#)
- [Virtualized NIC Requirements on page 35](#)

Overall Resource and Access Requirements

Ensure that the following resources are available:

- One or more vSphere ESX/ESXi 4.x hosts. Beginning with Release 5.0r2, vGW Series is enhanced to provide support for vSphere 5.0.

We recommend that you use more than one host for your deployment.



NOTE: vSphere 5.0 supports only ESXi hosts.

You use the VMware vSphere Client software to integrate the vGW Series with the VMware infrastructure.

- A VMware Virtual Center (vCenter) server, version 4.x. vGW Series 5.0r2 and later releases also support vCenter 5.0.

The vCenter VMware management server oversees the virtualization data center. The vCenter can be a physical server or a VM running on an MS Windows server.

The vCenter server automatically imports vGW Series components, and it adapts security as necessary when changes are made to the virtualized environment.

- Network connectivity.

The vGW Security Design VM must be accessible through HTTPS to allow access to the VMware Virtual Infrastructure API. Access to the VMware Virtual Infrastructure API is also required for autodiscovery of VM resources.

If you have access to the VMware Virtual Infrastructure API, you can connect a Web browser to the vCenter host (<https://vCenter-IP-address>).

- Domain Name System (DNS) and Network Time Protocol (NTP) services for some components.

The vGW Security VM requires NTP access to the center.

- One of the following supported Web browsers is required:
 - Microsoft Internet Explorer 7, 8, or 9
 - Mozilla Firefox 3 or later



NOTE: Localized (non-English) versions of browsers, such as the Japanese version of IE7, are not fully supported. However, most character sets including Japanese should display properly.

Virtual Appliance System Requirements

You can configure a network attached storage (NAS) device or a local datastore to use for both the vGW Security Design VM and the vGW Security VM. However, we recommend the following:

- Store the vGW Security Design VM on a NAS device so that it can be VMotioned.

You can allow the vGW Security Design VM to be migrated between hosts because it is not tied to a specific ESX/ESXi host.

Because the vGW Security Design VM can be sensitive to a slow NAS device, you should monitor it closely.

- Store each vGW Security VM on a local datastore on the ESX/ESXi host where it is installed.

A vGW Security VM is installed on and always remains associated with a single ESX/ESXi host. The vGW Security VM is configured not to be migrated through VMotion because it is specific to its host. If the host becomes unavailable, its vGW Security VM is not used for another host. It is applicable only to the host to which it was initially deployed. Allowing a vGW Security VM to migrate can cause problems in the virtualized environment.

If you must, you can place the vGW Security VMs on a NAS device. In this case, ensure that the vGW Security VMs are not VMotioned away from their designated hosts.



NOTE: Do not use a read-only datastore.

- The virtual appliances assume the following memory and require the following disk space:
 - vGW Security Design VM:
 - Memory: For vGW Series Release 5.5, you can configure up to 7 GB. For earlier versions, use 3 GB.
 - Disk space: 11 GB
 - vGW Security VM:
 - Memory: 512 MB



WARNING: Do not manipulate the memory size of the vGW Security VM. It must remain 512 MB.

- Disk space: 1.5 GB

vGW Series VMware vSwitch Requirements

VMware lets you create abstracted network devices called virtual switches (vSwitches). A vSwitch routes traffic internally between virtual machines, and it links to external networks. A vSwitch works somewhat like a physical Ethernet switch. It detects which virtual machines are logically connected to its virtual ports to enable it to forward traffic to the correct virtual machines.

A vSwitch can be connected to physical switches using physical Ethernet adapters, also referred to as uplink adapters, to join virtual networks with physical networks. This process is similar to connecting physical switches to create a larger network. Even though a vSwitch works like a physical switch, it does not have the advanced functionality of a physical switch.

For the vGW Series:

- You can map a vSwitch to one or more physical NICs on the ESX/ESXi host server, although this is not a requirement.
- You can configure features such as QoS traffic shaping on a vSwitch.

vGW Series interoperates with the following types of switches:

- Standard VMware Virtual Switch
- VMware Distributed Virtual Switch (DVS)
- Cisco Nexus 1000V device



WARNING: vGW Series creates a vmervice-switch for its own use. Do not make changes to its configuration or in anyway affect it.

VMware Port Group Requirements

In the VMware virtualized environment, port groups are used to aggregate multiple ports under a common configuration. They serve as an anchor point for virtual machines that connect to labeled networks. Each port group is identified by a network label. If port groups are configured, they are often mapped to VLANs, although this is not necessarily the case.

An administrator assigns to a port group a virtual network interface card (NIC) that connects a VM with a vSwitch.

There are two types of port groups:

- Virtual machine port groups. We recommend that you create a port group designated for a few test VMs to be secured.
- VM kernel port groups. This type of port group is used for storage for VMotion and ESX/ESXi host management.

We recommend that you create a port group designated for communication between the vGW Security Design VM and the management interfaces on each of the vGW Security VMs. For example, you might call this port group Juniper Networks vGW Management. You can associate this port group with a VLAN, but it must not filter TCP 443 or TCP 8443. There must be IP address space available for the vGW Security Design VM interface and for each of the vGW Security VMs.

You can use the preexisting VMware Management port group for this purpose.



WARNING: vGW Series creates two port groups that it uses to connect the VMsafe network to the vGW kernel module for communication. Do not change the configuration of these port groups in any way. They are for internal use only. The monitor port group is used when you activate console monitoring.

It is a promiscuous port group that vGW Series uses to view traffic into the service console for Network module tracking and IDS. You might notice that these port groups are created when you install vGW Series:

```
Altor_1VAm_Monitor_
vmservice-althor
vmservice-vmknic-pg
```

Virtualized NIC Requirements

Consider the following details when configuring vNICs:

- Do not change the vGW Security VM and the vGW Security Design VM vNICs default configuration. By default, vNICs are set to connect when they are powered on.
- VMs can use any type of vNIC supported by VMware. For example, their administrators might want to use VMXNET3 for improved performance. Additionally, the vGW Series supports multiple vNICs for a VM. You can secure these vNICs with different policies using the Policy per vNIC feature. For details, see *Configuring the vGW Series Policy per vNIC Feature*.

Related Documentation

- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 40](#)
- [Understanding vGW Series on page 3](#)
- [Integrating the vGW Series with VMware Using the Settings Module](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts](#)

Preparing to Integrate the vGW Series with the VMware Environment

Before you import and install the vGW Series with the VMware environment, take the following precautions:

- Ensure that the virtual machines (VMs) in your environment can communicate with one another. You can use the ping command for this purpose.
- Ensure that you can access the ESX/ESXi hosts using SSH.
- Verify that Domain Name System (DNS) and Network Time Protocol (NTP) services are functioning properly on the network.
- Ensure that there is HTTP access to the datastore. By default, vCenter allows this access. However, if you have hardened your configuration by adding false to your vpxd.cfg file, vGW Series will not be able to automatically deploy the vGW kernel module.

Related Documentation

- [Understanding the VMware Infrastructure and vGW Series on page 7](#)
- [Understanding vGW Series on page 3](#)

- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)

Understanding vGW Series Environment Time Synchronization

vGW Series uses NTP to synchronize all times in the environment to ensure that all security policies, logs, and other time-based data, are properly marked. Large time gaps resulting from improperly configured systems can cause unexpected results and make it difficult to troubleshoot the environment.



WARNING: Do not edit the time settings, including the time zone, for a vGW Security Design VM directly. The vGW Security VMs obtain their time setting from the vGW Security Design VM. Inappropriate changes made to the vGW Security Design VM can adversely affect all vGW Security VMs.

Take care in configuring NTP when you install vGW Series. Ensure that the internal NTP server is up and functioning properly, and that outbound Internet access to an Internet time server is available.

See *Configuring vGW Series Time Settings*.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 63](#)
- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability](#)
- [Installing a Secondary vGW Security VM for High Availability](#)

vGW Series VMsafe Firewall + Monitoring and VMsafe Monitoring Modes

vGW Series supports two modes that allow you to secure and observe virtual machines (VMs) and traffic:

- VMsafe Firewall + Monitoring mode allows you to secure specific VMs or entire port groups. It provides visibility into all traffic that transits each protected VM.
- VMsafe Monitoring mode allows you to fully monitor VMs without securing them. It guarantees that no packets are blocked because of an incorrectly configured security policy. When monitoring mode is enabled, security policies are not loaded into the vGW hypervisor kernel module.

When this option is selected and you create a group for which you do not select the Policy Group option, the group is automatically placed in the Monitoring Groups section of the VM tree. (To create a group, use the Settings module Security Settings > Groups page.) For details on creating groups, see *Understanding vGW Series Groups*.

To use VMsafe Monitoring mode:

1. First enable it. On the Settings module vGW Application Settings > Install Settings page, select the **Enable Monitoring-only option for VMsafe** checkbox.
2. If VMsafe Monitoring mode is enabled, when you use the Settings module vGW Application Settings > Installation page to install a vGW Security VM, you can choose either VMsafe Firewall+Monitoring or VMsafe Monitoring for the installation.

To view monitoring groups, use the Settings module > Security Settings > Groups page, and select **Monitor Groups**. The table shows all configured monitoring groups.

When vGW Series installs a vGW Security VM on an ESX/ESXi host in either VMsafe Firewall + Monitoring mode or VMsafe Monitoring mode, it uses the VMware VMsafe networking APIs to build the security engine as a kernel module in the hypervisor of the host.

This installation allows for full protocol inspection of every VM.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 64](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts](#)
- [Understanding the vGW Security Design VM on page 63](#)

CHAPTER 5

OVA and vGW Series Deployment

- [Understanding the Open Virtualization Format OVA Template Method on page 39](#)
- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 40](#)
- [Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware on page 49](#)
- [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 51](#)

Understanding the Open Virtualization Format OVA Template Method

vGW Series leverages the Open Virtualization Format (OVF) standard for packaging and delivering virtual machines (VMs). The OVF supports industry-standard content verification and integrity checking, and it provides a basic scheme for managing software licensing. As described by the standard, OVF defines an "open, secure, portable, efficient and extensible format for the packaging and distribution of software to be run in virtual machines." The standard also supports the OVA template method of packaging and distributing software in a single archive. You use the vSphere 4.x client to load the OVA file.

You can use OVA to deploy the vGW Series VMs in the following ways:

- In a single bundled OVA package, also referred to as a Combo Package, that contains the vGW Security Design VM and the vGW Security VM template.

vGW Series uses OVA to deliver a single file containing both vGW Series components.

- In nonbundled OVA files to separately deploy the vGW Security Design VM and the vGW Security VM template.

The OVA Combo Package installs a vApp, and VMware will not install a vApp on a cluster for which the Dynamic Resource Scheduler (DRS) is not enabled. You can take the nonbundled approach in this case.

The nonbundled OVA approach is also useful for installing the most current vGW Security VM, after the initial installation, to ensure that the latest version is used for automatic vGW Security VM instantiation on ESX/ESXi hosts.

You can also use it to create a secondary vGW Security Design VM for high availability. For details, see *Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability*.

Related Documentation

- [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 51](#)
- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 40](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 35](#)
- [Understanding the VMware Infrastructure and vGW Series on page 7](#)
- [Understanding vGW Series on page 3](#)

Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure

This topic explains how to integrate the vGW Series appliances—the vGW Security Design VM and the vGW Security VM template—with the VMware virtualized infrastructure.

For information on “vGW Series Prerequisites and Resource Requirements for the VMware Environment” on page 31.

This topic includes the following sections:

- [Requirements on page 40](#)
- [Overview on page 40](#)
- [Downloading the vGW Series OVA Combo Package on page 41](#)
- [Integrating the vGW Series with the VMware Infrastructure on page 41](#)

Requirements

For information, see “vGW Series Prerequisites and Resource Requirements for the VMware Environment” on page 31.

Overview

The bundled OVA template allows you to deploy both the vGW Security Design VM and the vGW Security VM appliances in a single OVA archive file. In this case, OVA creates a single vApp and inserts the two vGW Series appliances into it.

You can delete the vApp after the deployment and integration process is complete. It is used only to convey the vGW Series VMs. However, take care not to delete it before then.

In the single Combo Package file, the OVA template deploys:

- The vGW Security Design VM
- The vGW Security VM

You must manually convert the vGW Security VM to a template after you integrate the vGW Series with the VMware infrastructure. vGW Series uses the resulting template to instantiate a vGW Security VM on each ESX/ESXi host when you secure that host.



NOTE: The OVA Combo Package installs a vApp, and VMware will not install a vApp on a cluster for which DNS is not enabled. In this case, you must use the nonbundled OVA method to deploy each component separately.

- For details, see [“Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware”](#) on page 49.
- For details, see [“Using the OVA Single File Method to Integrate the vGW Security VM with VMware”](#) on page 51.

Downloading the vGW Series OVA Combo Package

- Step-by-Step Procedure** To download the Juniper Networks OVA archive file that contains both the vGW Security Design VM and the vGW Security VM:
1. Navigate to the Juniper Networks Support page.
 2. Select **Software Downloads** from the Support box in the left column.
 3. Select **vGW (Altor)** in the Security pane.
 4. Select the **Software** tab.
 5. Click **vGW Series 5.5 Combo Package**, and log in to the site to download the file.

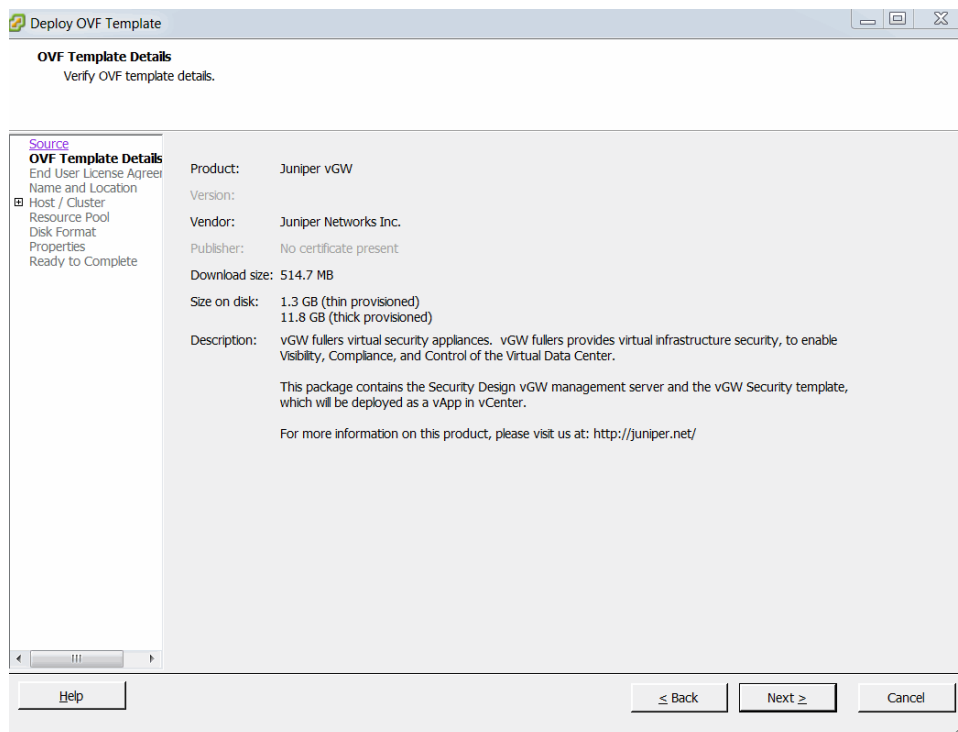
Integrating the vGW Series with the VMware Infrastructure

- Step-by-Step Procedure** To deploy the vGW Series appliances—the vGW Security Design VM and the vGW Security VM—and integrate them with the VMware infrastructure:
1. Using the vSphere client, load the bundled OVA file. Select **Deploy OVF Template** from the File menu.
 2. Enter the download filename or its URL in the Deploy from file or URL box—for example, enter: `c:\temp\vGW_Combo_5.5_#-#-#_#-#-#.ova`—and click **Next**.

You use the OVF template method to deploy the OVA file. After you specify the name of the OVA file and its location, the Appliance Wizard displays the OVA template details dialog box.
 3. Verify the contents of the OVA package, and click **Next**.

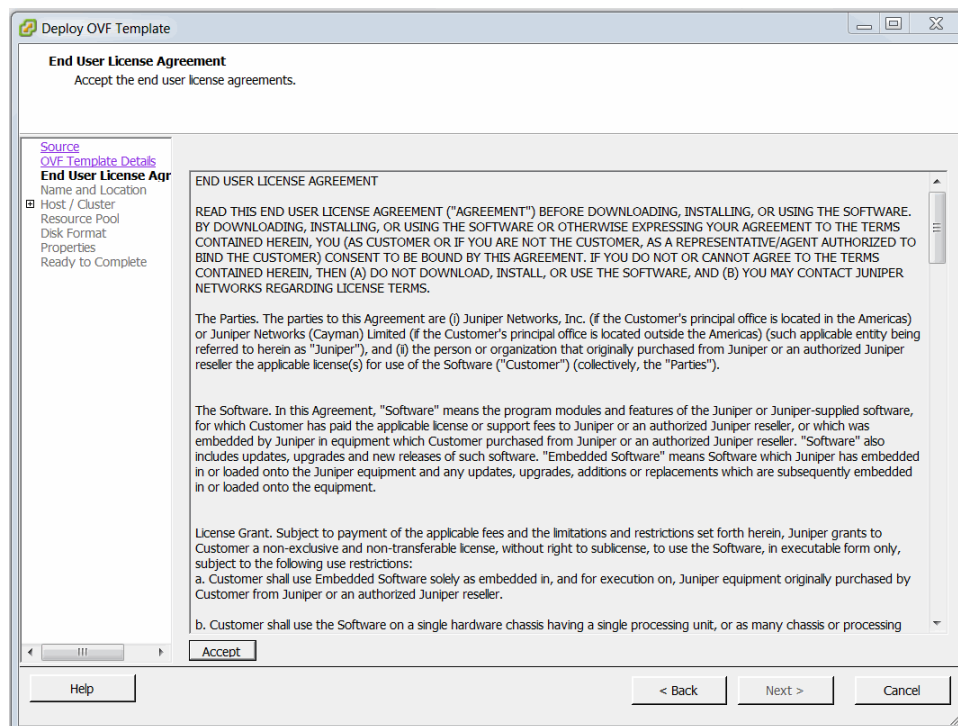
Before the wizard unbundles the OVA package, verify that it contains the vGW Series appliances. The OVA template summary also specifies the disk space requirements for thick and thin provisioning. See [Figure 17 on page 42](#).

Figure 17: OVA Template Details Page



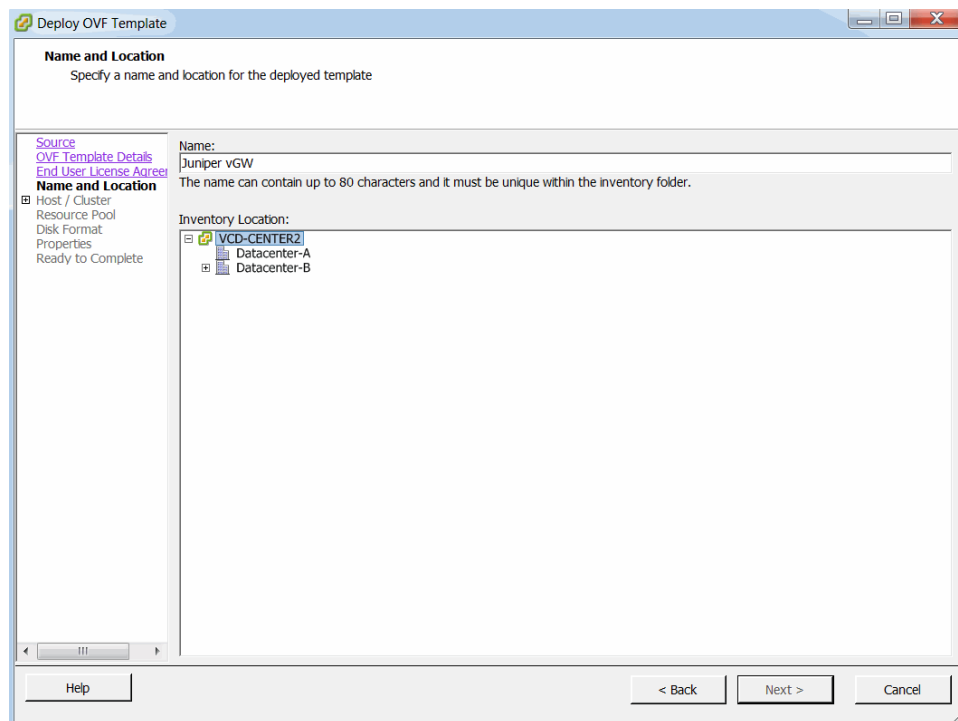
4. Accept the vGW Series license agreement, and click **Next**. See Figure 18 on page 42.

Figure 18: OVA File Deployment License Agreement



5. Specify a name for the vApp that will be created and a storage location. See [Figure 19 on page 43](#).

Figure 19: Naming the vApp



6. Specify the host or host/cluster on which to run the deployed template. We recommend that you use a network storage device (NAS) so that it can be migrated through VMotion for space optimization. See [Figure 20 on page 44](#).

Figure 20: Specifying the Host and Cluster

Deploy OVF Template

Host / Cluster
On which host or cluster do you want to run the deployed template?

Source
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
Host / Cluster
 Specific Host
 Resource Pool
 Disk Format
 Properties
 Ready to Complete

Datacenter-B
 10.159.24.176
 10.159.24.183

Help < Back Next > Cancel

7. Select the datastore. Do not use a read-only datastore. [Figure 21 on page 44.](#)

Figure 21: Selecting the Storage

Deploy OVF Template

Storage
Where do you want to store the virtual machine files?

Source
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
[Host / Cluster](#)
Storage
 Disk Format
 Network Mapping
 Properties
 Ready to Complete

Select a destination storage for the virtual machine files:
 VM Storage Profile: Select...

Name	Drive Type	Capacity	Provisioned	Free	Type	Storage	Thin Provisioning	Access
10.159.24.3	Unknown	366.76 GB	257.00 ...	109.77 GB	NFS		Supported	Single ho...
Local183	Unknown	5.00 GB	393.00...	4.62 GB	VMFS3		Supported	Single ho...
SLT-NAS	Unknown	5.40 TB	1.92 TB	3.49 TB	NFS		Supported	Multiple ...

☐ Disable Storage DRS for this virtual machine

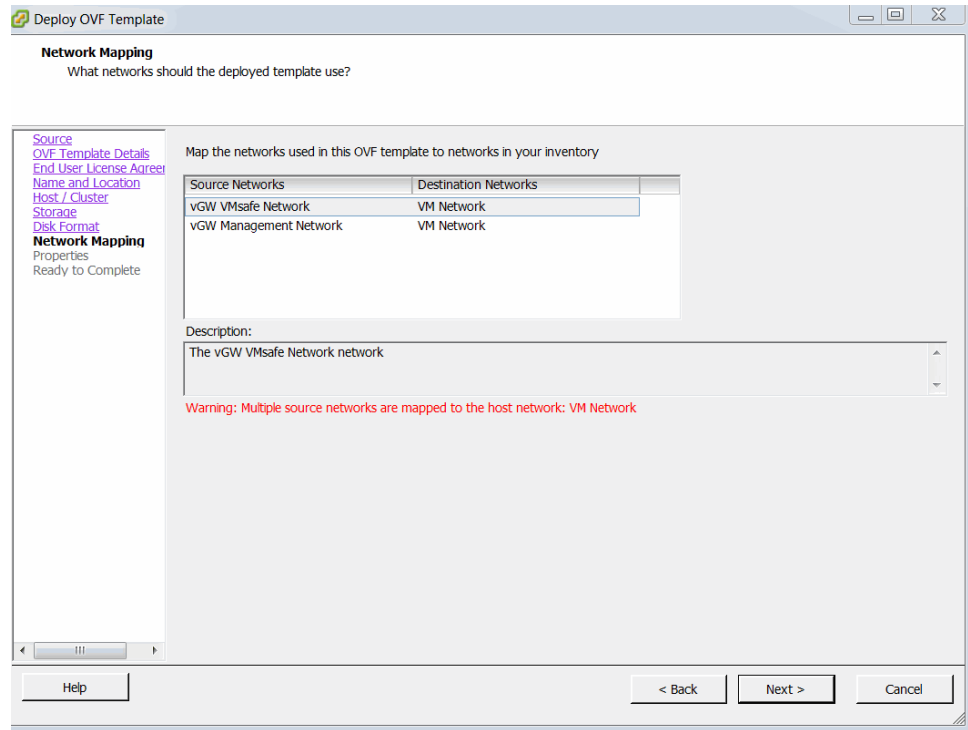
Select a datastore:

Name	Drive Ty...	Capacity	Provisioned	Free	Type	Thin Provisioning	Access

Help < Back Next > Cancel

8. Select the disk format. Accept the thick provisioned format default. Thick provisioning preallocates all required space for the product.
9. Map the networks. Set the vGW management network to a destination network that is accessible to vCenter and the vGW Security Design VM. See [Figure 22 on page 45](#).

Figure 22: Mapping the vGW Management Networks



10. Specify the size of the database to use for storing vGW Series files.
The database stores network connection records and firewall logs.
See [Figure 23 on page 46](#).

Figure 23: Specifying the Database Disk Size

Deploy OVF Template

Properties
Customize the software solution for this deployment.

Source
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
[Host / Cluster](#)
[Storage](#)
[Disk Format](#)
[Network Mapping](#)
Properties
 Ready to Complete

Database Disk size
 Please choose the size (in GB) of the virtual disk to be created for the database used for network session and firewall log data.
 Minimum Size = 2GB
 Recommended size = 8GB

8

Help < Back Next > Cancel

The default disk size is 8.0 GB. In a typical environment that includes 5 to 10 ESX/ESXi hosts, a database of this size can accommodate data accumulated over several months. However, for your environment you might want to deploy a database that is larger than 8.0 GB.

You can increase the database size later if you find that the current space is not adequate. Although there is no hard-coded limit, we recommend restricting the size to less than 75 GB.

11. Verify that the configuration is correct, and click **Finish** to complete the deployment. See [Figure 24 on page 46](#).

Figure 24: Verifying That the Configuration Is Correct

Deploy OVF Template

Ready to Complete
Are these the options you want to use?

Source
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
[Host / Cluster](#)
[Storage](#)
[Disk Format](#)
[Network Mapping](#)
[Properties](#)
Ready to Complete

When you click Finish, the deployment task will be started.

Deployment settings:

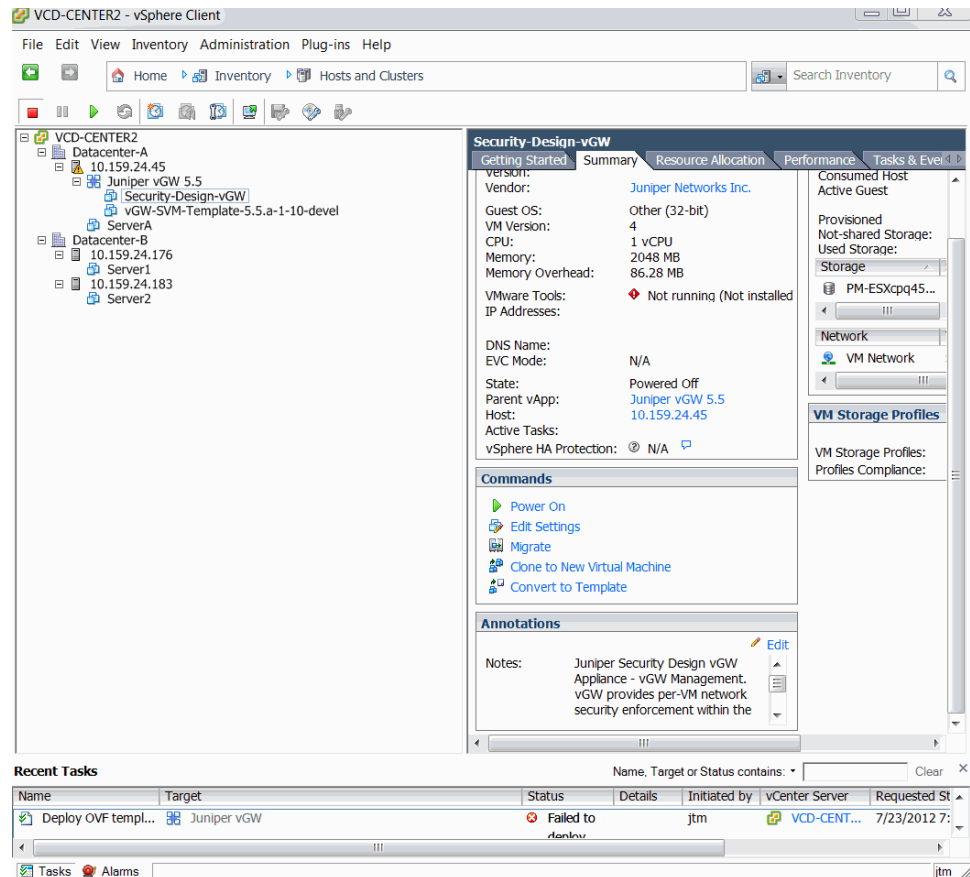
OVF file:	C:\BACKUP\ALTOR\Altior 6.0, 2012, Q1\CVBCs for 6.0\vGW_Combo_5.5_c-2-7_a-1-10_devel.ovf
Download size:	514.7 MB
Size on disk:	1.3 GB
Name:	Juniper vGW
Folder:	Datacenter-B
Host/Cluster:	10.159.24.183
Datastore:	SLT-NAS

The Virtual Appliance Wizard downloads the files and inserts the vGW Series VMs as a single virtual appliance (vApp) into the VMware infrastructure.

When the OVA import is completed, the vCenter includes the vApp containing both the vGW Security Design VM and the vGW Security VM template components.

12. Expand the appliance called Juniper vGW 5.5 to display the vGW Security Design VM and the vGW Security VM. See [Figure 25 on page 47](#).

Figure 25: Displaying the vGW Appliance Components



Move the two vGW Series VMs out from the vAPP. Afterward you can delete the vApp if you choose to, but it is not necessary.



NOTE: vGW Series uses the VMware vApp deployment feature as a vehicle to deliver multiple VMs in the same OVA file. The vApp structure is redundant after it is used for deployment, and therefore you can delete it. However, do not delete the vApp without first having moved the vGW Series VMs out from it. If you do, the newly created vGW Series VMs would be deleted when you delete the vApp.

After you remove the vGW Series VMs from the vApp:

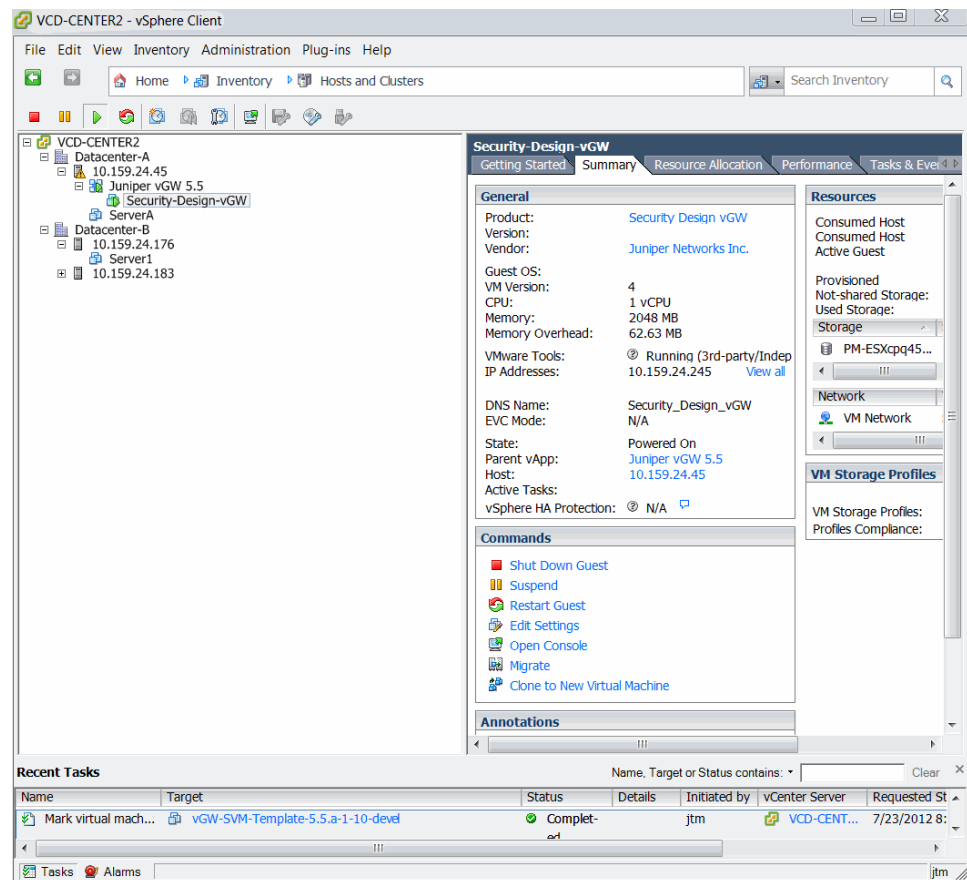
- a. Convert the vGW-SVM-Template VM to a template that the vGW Security Design VM and installer can use to instantiate a vGW Security VM on each ESX/ESXi host to be secured.

Right-click the template, select **Template**, and select **Convert to Template**.

- b. Right-click the vGW Security Design VM and power it on. [Figure 26 on page 49](#) shows the vCenter summary information for the vGW Security Design VM.

[Figure 26 on page 49](#) shows the vCenter summary information for the vGW Security Design VM.

Figure 26: vGW Security Design VM Summary Tab in vCenter



Related Documentation

- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the VMware Infrastructure and vGW Series on page 7](#)
- [Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware on page 49](#)
- [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 51](#)
- [Understanding the vGW Security Design VM on page 63](#)
- [Understanding the vGW Security VM on page 64](#)

Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware

This topic explains how to download and deploy a single OVA file containing the vGW Security Design VM.

To download an OVA file containing the vGW Security Design VM appliance and deploy it:

1. Download the Juniper Networks vGW OVA file.
 - a. Navigate to the Juniper Networks Support page.
 - b. Select **Software Downloads** from the Support box in the left column.
 - c. Select **vGW (Altor)** in the Security pane.
 - d. Select the **Software** tab.
 - e. Click **Security Design vGW 5.5**, and log in to the site to download the file.
2. Load the OVA file for the vGW Security Design VM using the vSphere 4.x client (File > Deploy OVF Template), and enter the name of the OVA download file in the Deploy from file or URL box.

For example, enter: `c:\temp\SecurityDesignvGW.ova`.

3. Follow the Virtual Appliance Wizard process, and select the appropriate options for your environment.
4. Click **Finish** to download the files and integrate the vGW Security Design VM with the VMware infrastructure.

After the vGW Security Design VM import process is completed, you must add a virtual hard disk for it.



NOTE: This step is not required for the bundled approach because it is done automatically.

The default disk size is 8.0 GB. In a typical environment that includes 5 to 10 ESX/ESXi hosts, a database of this size can accommodate data accumulated over several months.

For your environment you might want to deploy a database larger than 8.0 GB. Note that you can increase the database size later if you find that the current space is not adequate. The disk should not be thin-provisioned.

5. Add a disk to be used as the datastore:
 - a. Select **vGW Security Design VM**.
 - b. Select the **Summary** tab, and click **Edit Settings > Add a Hard Disk virtual device**.

This disk is used for the database that stores network connection records and firewall logs. Select a NAS device so that VMotion can be used to migrate the datastore.

6. Power on the vGW Security Design VM.

- Related Documentation**
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)
 - [Understanding vGW Series on page 3](#)
 - [Understanding the VMware Infrastructure and vGW Series on page 7](#)
 - [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 51](#)
 - [Understanding the vGW Security Design VM on page 63](#)
 - [Understanding the vGW Security VM on page 64](#)

Using the OVA Single File Method to Integrate the vGW Security VM with VMware

To download a nonbundled OVA file containing the vGW Security VM and deploy it:

1. Load the OVA file for the vGW Security VM using the VMware vSphere Client (File > Deploy OVF Template), and insert the template name vGW-SVM-Template in the Deploy from file or URL box. For example, enter **c:\temp\vGW-SVM-Template.ova**.
2. Select the appropriate options for your environment in each of the steps presented by the Virtual Appliance Wizard.

Configure the host/cluster, resource pool, and so on, that is appropriate for your environment.

When you are asked for network mapping information, accept the default settings. vGW Series automatically configures these settings later.

3. When the Virtual Appliance Wizard completes, right-click the resulting VM and select **Template > Convert to Template**.

You can use the resulting template to automate installation of vGW Security VMs on ESX/ESXi hosts to secure parts of your virtual network. The vGW Security Design VM and installer require the template to instantiate the vGW Security VM on hosts to be secured.

- Related Documentation**
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 31](#)
 - [Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware on page 49.](#)
 - [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 40](#)
 - [Preparing to Integrate the vGW Series with the VMware Environment on page 35](#)

PART 3

vGW Series Setup

- [SDVM Set Up Process on page 55](#)

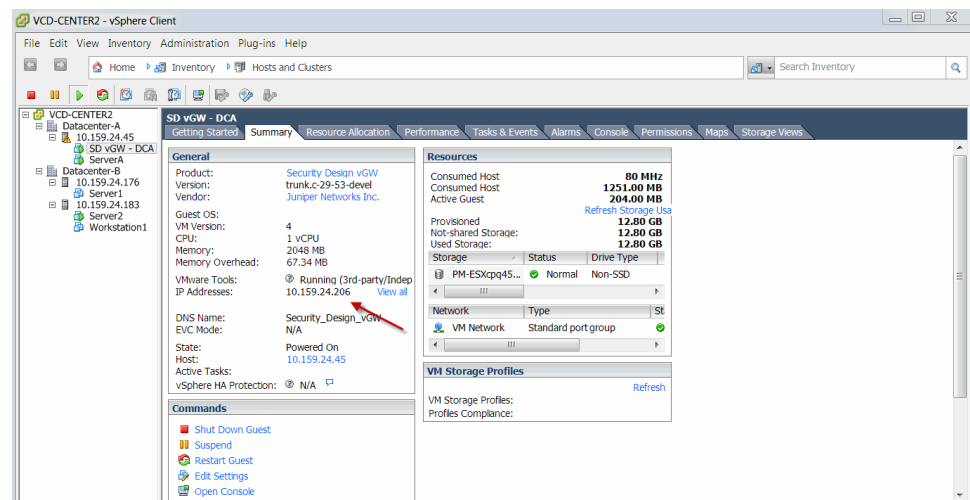
CHAPTER 6

SDVM Set Up Process

- Setting Up vGW Series on page 55

Setting Up vGW Series

After you download and deploy the vGW Series and power on the vGW Security Design VM, you can configure basic operating system parameters such as the vGW Security Design VM IP address. By default, the vGW Security Design VM is configured for You use a Web browser to access the vGW Security Design VM.



If DHCP is available, you can determine the IP address from the vCenter server. To do so, select the **vGW Security Design VM** in the vCenter console, and then select the **Summary** tab. Alternatively, you can display the IP address by selecting the **Console** tab.

By default, the vGW Security Design VM is configured for dual stack, with IPv4 configured to use DHCP and IPv6 configured to use stateless autoconfiguration.

The IPv6 address can be configured in the following ways, depending on the netirj

- Stateless IPv6 network configuration—If your network includes IPv6-enabled routers, the vGW Security Design VM automatically configures itself with a valid IPv6 address that is reachable from machines within your network and outside it.

- Link-local IPv6 address—If your network does not include IPv6-enabled routers, the vGW Security Design VM configures itself automatically with an IPv6 link-local address. Link-local addresses can be used to reach nodes attached to the same link, and therefore there is no access outside the local network.

For IPv4, if DHCP is not available on the vGW Security Design VM network:

1. Log in to the console using **admin** for both the username and the password.
2. At the command prompt, enter **config network**, and specify the options to assign an IP address.

After an IP address is set through DHCP or a static IP address is assigned, you can access the vGW Security Design VM using a Web browser.

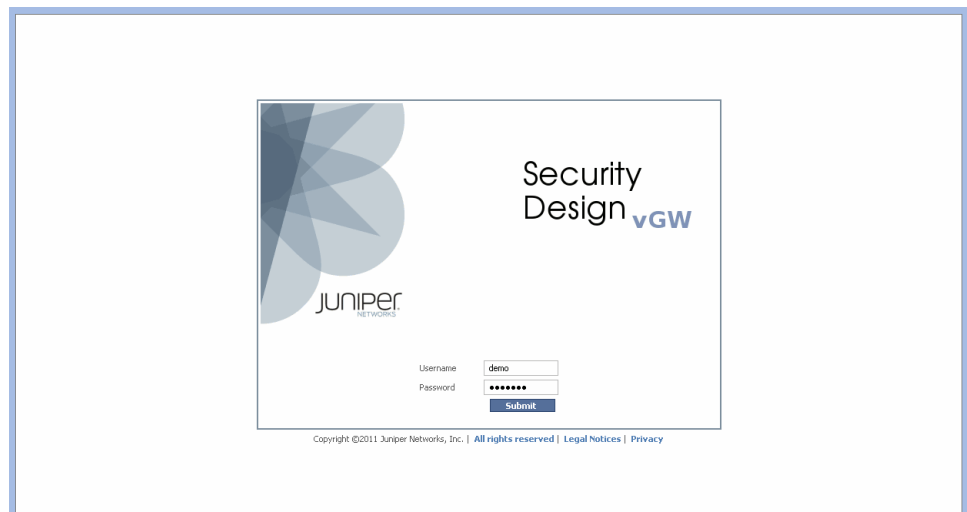
3. Using a supported Web browser, connect to the vGW Security Design VM management interface through HTTPS.

Enter **admin** for both the username and password. See [Figure 27 on page 56](#).

vGW Series supports the following Web browsers:

- Microsoft Internet Explorer 7 and 8
- Mozilla Firefox 3 or later

Figure 27: vGW Series Security Design VM Login Screen



4. Read the information message, and review the process overview shown in the Wizard Progress pane.
5. Change the default vGW Global Admin account password—admin—that you used to log in.

You must change the default password. Store the new password in a secure location. It is difficult to recover a lost or forgotten password.



TIP: You can integrate administration accounts with the vGW Security Design VM after the installation is complete.

6. Configure networking parameters for the vGW Security Design VM.

Set the correct destination network for vGW Management Network and leave the VMsafe Network unchanged.

7. If you changed the IP address, you must log in to the system again. Changes to the IP address take effect immediately.



NOTE: The system tries to verify the DNS server entries. You can safely ignore the warning message.

8. Set the system time.

Set the correct time zone, and then specify the NTP servers for your environment.

The vGW Series components require that the correct system time be set on all ESX/ESXi hosts.

If you do not have an NTP server, you can use a predefined server. If you do not have outbound Internet access to contact the NTP servers and if you do not have an internal NTP server, then you must clear all entries shown in this window and set the time manually.

To set the time manually, log in to the console, and use the vGW Series command-line utility.

At this point, the wizard confirms that a database disk was created and initialized properly. If you have not defined the database disk properly, the wizard displays a message.

If you are using a 30-day evaluation license, you can continue to use the vGW Series in that mode, or you can enter your permanent licenses. See [Figure 28 on page 58](#).

Figure 28: vGW Series Licensing

The screenshot shows the 'vGW Installation Wizard' window with the 'Product Licensing' step selected in the 'Wizard Progress' sidebar. The main content area contains the following text:

Product Licensing

Enter permanent license or continue in evaluation mode.

You can select to continue in a full function evaluation mode for 30 days or input your permanent license(s) now. If you continue in evaluation mode at the end of the trial period no locking of traffic occurs. However, features won't be usable and you will need to enter a new license or uninstall the product.

To get a permanent license or a long term evaluation license, please contact your Juniper Networks reseller or Sales.

Enter Permanent or long term evaluation license(s)

Buttons: **Prev Step** | **Continue in 30 Day Eval. Mode**

9. Insert the appropriate authentication information, then click **Next** to perform an authentication test. See [Figure 29 on page 58](#).

Figure 29: Authenticating to the vGW Installation Wizard

The screenshot shows the 'vGW Installation Wizard' window with the 'vCenter Settings' step selected in the 'Wizard Progress' sidebar. The main content area contains the following text:

vCenter Settings

Establish connection to vCenter: [more](#)

Server Name or IP Address:

Username:

Password:

Select a scope for your Security Design vGW: [more](#)

☒ Entire vCenter ☐ Selected Datacenters

Buttons: **Prev Step** | **Next Step**

For the vGW Series to query the vCenter for the VM inventory and other operations, you must have an account with read and write access.

- If the connection works properly, a message appears that shows the number of ESX/ESXi hosts and VMs that were discovered.

- If there is a connection issue, you are notified. In that case, ensure that you have the correct credentials and that IP connectivity to the vCenter system exists.

In some cases, you may need to insert another vNIC into the vGW Security Design VM. Under these circumstances, you must connect that vNIC to the network that connects to the vCenter server.

10. (Optional) Configure the e-mail server that you want to use to send reports.

Using this option, you can configure the vGW Series to send reports on system activity through e-mail. Additionally, you can configure basic information used in the report, such as the subject, the content of standard report e-mail, and so on. After you configure these parameters, you can test the e-mail connection.

You can also configure this information later or change it after the installation completes using vGW Security Design VM > Settings > Applications section.

11. Define a template for deploying the vGW Security VMs to secure the environment.

If you have not downloaded the vGW Security VM and converted it to a template, do so now. You can define:

- How the vGW Series responds when a VM tries to connect to an ESX/ESXi host on which the VMsafe kernel module cannot be loaded or is not present.
- Whether the VMsafe Monitor Mode installation page appears.

Unless you plan to deploy the product in monitor mode, leave the Monitoring-only option for VMsafe unchecked.

Also, unless you want to drop network traffic to VMs when the vGW Series fails to load, you should leave the default option of Allow All traffic. You can change this option later, if you want to change the behavior for one or more VMs.

12. Click **Done** to complete the vGW Security Design VM setup.

The vGW Security Design VM appears. You use this module to deploy vGW Security VMs to the ESX/ESXi hosts to be secured, to configure other vGW Series features, and to view specific and summary results information and reports.

Related Documentation

- [Preparing to Integrate the vGW Series with the VMware Environment on page 35](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module](#)
- [Understanding the vGW Series Main Module on page 16](#)
- [Understanding Licenses for the vGW Series](#)

PART 4

SDVM and SVM

- [Basics on page 63](#)
- [SDVM Navigation and VM Tree on page 67](#)

CHAPTER 7

Basics

- [Understanding the vGW Security Design VM on page 63](#)
- [Understanding the vGW Security VM on page 64](#)
- [Understanding the vGW Series Kernel Module on page 64](#)

Understanding the vGW Security Design VM

vGW Series includes a management center called the vGW Security Design VM. The vGW Security Design VM allows you to create multi-tiered policies to protect and secure VMs. You use it to push those policies to vGW Security VMs which are installed on the hosts that they secure.

The vGW Security Design VM modules provide features that allow you to perform the following tasks and many others:

- Perform network traffic analysis.
- Configure your environment to protect against intrusions and attacks.
- Quarantine suspect files and VMs.
- Inspect for malware and anomalous behavior.
- Configure and generate reports providing information on all aspects of your monitored and secured environment.
- Create dynamic groups called Smart Groups that secure VMs automatically in varying ways based on characteristics of the VM without your needing to intervene.

You use the vGW Security Design VM to configure vGW Series. You must use a Web interface browser to access the vGW Security Design VM by its IP address. You can obtain the IP address by clicking the **Summary** tab for the Security Design VM in VMware vCenter.

After you initially bring up the vGW Security Design VM, you can access it by entering **admin** for the username and entering the password that was set during installation. To log out of the vGW Security Design VM, click **logout** in the upper right corner of the vGW Security Design VM page.

When you log in to the vGW Security Design VM, you see the Main module page with its Dashboard tab. See [Figure 30 on page 64](#). The page displays information gathered from the activity of various vGW Security Design VM modules.

Figure 30: Main Module Displayed at Login



Related Documentation

- [Understanding the vGW Security Design VM Taskbar on page 69](#)
- [About the vGW Security Design VM Tree on page 70](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding vGW Series on page 3](#)

Understanding the vGW Security VM

To secure the virtual machines (VMs) on an ESX/ESXi host, you install a vGW Security VM on the host. In turn, the vGW Security VM installs the vGW kernel module into the VMware hypervisor of the host. The vGW Security VM acts as a conduit between the vGW Security Design VM and the vGW kernel module. The vGW Security VM also maintains policy and logging information.

You use the vGW Security Design VM to deploy vGW Security VMs to ESX/ESXi hosts. You use it to create firewall policies that are pushed to the vGW Security VMs. The vGW Security VM inserts the policies into the vGW kernel module where all connection enforcement occurs.

Related Documentation

- [Understanding the vGW Security VM Settings](#)
- [Understanding the vGW Series Firewall Module](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts](#)
- [Installing a Secondary vGW Security VM for High Availability](#)

Understanding the vGW Series Kernel Module

vGW Series kernel module is the policy enforcement engine that is loaded into the hypervisor of an ESX/ESXi host to be secured. It utilizes the VMware VMID to ensure that

the correct policy is applied to a VM. It manages state synchronization in order to support VMotion.

It is a lightweight component that plugs directly into the host's hypervisor—without relying on an OS or a VM.

Communication between the vGW Series kernel module in the ESX/ESXI host's hypervisor and the vGW Security VM occurs over a special VMware vmservice vSwitch.

vGW Security VM is the conduit to the vGW Series kernel module. It inserts security policy into the kernel module, transfers logs and network information from the kernel module to the vGW Security Design VM and other devices SYSLOG, NetFlow V9 devices.

**Related
Documentation**

- [Understanding the vGW Security VM on page 64](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 63](#)

CHAPTER 8

SDVM Navigation and VM Tree

- Understanding vGW Security Design VM Navigation on page 67
- Understanding the vGW Security Design VM Taskbar on page 69
- About the vGW Security Design VM Tree on page 70

Understanding vGW Security Design VM Navigation

You use the vGW Security Design VM taskbar in conjunction with the VM tree to navigate the graphical user interface. The combination of the two allows you to select VMs and view and configure information about selected ones. See [Figure 31 on page 67](#) and [Figure 32 on page 68](#).

You can use the search field to filter VMs in various ways. See “[About the vGW Security Design VM Tree](#)” on page 70.

Figure 31: vGW Security Design VM Taskbar





Figure 32: VM Tree




Related Documentation

- [Understanding the vGW Security Design VM Taskbar on page 69](#)
- [About the vGW Security Design VM Tree on page 70](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 63](#)

Understanding the vGW Security Design VM Taskbar

The taskbar lets you select the vGW Security Design VM module to use and move from one module to another. The vGW Security Design VM provides a modular configuration and information display structure. The taskbar includes icons representing the various modules. [Figure 33 on page 69](#) shows the taskbar.

Figure 33: vGW Security Design VM Taskbar



[Table 4 on page 69](#) identifies the modules that the vGW Security Design VM icons represent. You can click the link for a module to go to a topic that covers it.

Table 4: Taskbar Icons








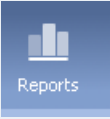

Icon	Module	Description	For details, see:
	Main	Combines status, alerts, and network activity into a single view. It identifies the VMs that are quarantined and allows you to take action on them.	“Understanding the vGW Series Main Module” on page 16
	Network	Displays a network activity summary, top protocols, sources, destinations, talkers, and connections.	Understanding the vGW Series Network Module
	Firewall	Manages and installs policies, and displays logs.	Understanding the vGW Series Firewall Module
	IDS	Monitors all network traffic or a selected subset of VMs or protocols.	Understanding the vGW Series IDS Module
	AntiVirus	Protects VMs by detecting malware, and it quarantines affected files or VMs.	Understanding vGW Series AntiVirus

Table 4: Taskbar Icons (*continued*)

Icon	Module	Description	For details, see:
 Introspection	Introspection	Scans systems and reports on the software running in each VM (operating systems, patch-levels, and applications). It includes an Image Enforcer feature that allows you to specify VM templates or active VMs whose configurations are used as Gold Image comparison points. Contents of VMs are compared against the Gold Image.	<i>Understanding the vGW Series Introspection Module</i>
 Compliance	Compliance	Monitors the virtual infrastructure against a predefined set of rules to guarantee all components are configured securely.	<i>Understanding the vGW Series Compliance Module.</i>
 Reports	Reports	Produces detailed system and security reports.	<i>Understanding the vGW Series Reports Module.</i>
 Settings	Settings	Controls configuration settings, including passwords.	<i>Understanding the vGW Series Settings Module</i>

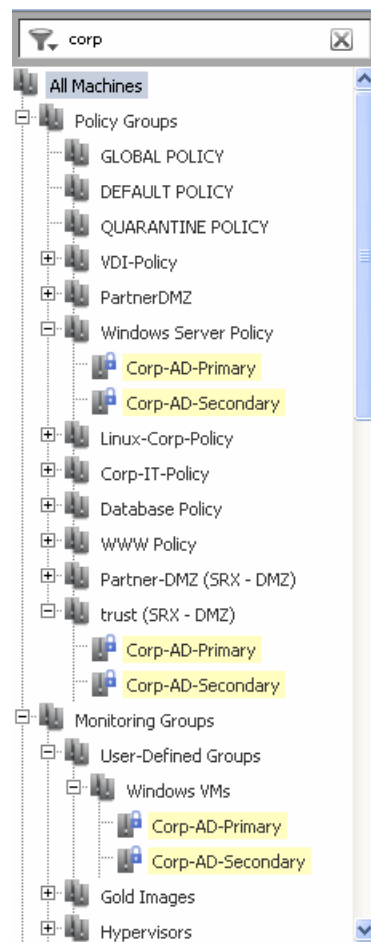
- Related Documentation**
- [About the vGW Security Design VM Tree on page 70](#)
 - [Understanding the vGW Security Design VM on page 63](#)
 - [Understanding vGW Series on page 3](#)

About the vGW Security Design VM Tree

You use the VM tree in conjunction with the vGW Security Design VM modules. The VM tree lets you select virtual machines (VMs) to focus on, configure, and view information about.

You can select a group of VMs or an individual VM in the VM tree either by clicking its name or using the filter box. [Figure 34 on page 71](#) shows VMs belonging to three groups.

Figure 34: VM Tree with Selected VMs



See “Understanding the vGW Security Design VM Taskbar” on page 69.

- [VM Tree Overview on page 71](#)
- [Locating VMs in a Complex VM Tree on page 72](#)

VM Tree Overview

In conjunction with the selected vGW Security Design VM module, the VM tree controls the information displayed in the pane beside it. You can select all VMs in the tree, groups of VMs, or a single VM. When you select a module using the taskbar, that module's content appears as it applies to the VMs that you selected in the VM tree. The module controls the type of information that appears; the tree controls the VMs whose information appears. The combined selections allow you to configure or view information for that module as it pertains to the VMs. For example, to view network traffic for all machines, select **All Machines** in the tree, and then click the Network icon in the taskbar.

The VM tree contains the following main groups:

- Policy Groups

Contains all security policy groups, including Global, Default, and Quarantine. It also contains Illegal IPv4 Sources and Illegal IPv6 Sources groups and any policy groups that you define.

- Monitoring Groups

Contains all groups that were created with the Policy Group option, groups for monitoring the Hypervisor and Compliance state, and a group containing VMs or templates used as Gold Images by the Introspection module's Image Enforcer feature.

- Monitored/Secured VMs







Lists VMs monitored by the vGW Series, VMs that have a firewall protecting their network traffic, or both.

- Unsecured Machines

Lists all VMs that are not currently being analyzed or protected by the vGW Series.

[Table 5 on page 72](#) identifies the icons that show the state of monitored VMs.

Table 5: Virtual Machine State Icons

	The VM is being fully monitored, but it is not secured. For example, no firewall policy is loaded.
	The VM or the externally defined machine is not being monitored, and it has not been moved to a network secured by vGW Series. NOTE: Network reports can display sessions between an unmonitored system and a monitored VM.
	vGW Series cannot determine the IP address of the machine. This could be because it is powered down, suspended, or does not have VMware Tools installed. TIP: You can manually define an IP address by selecting the Settings module's vGW Application Settings > Machines .
	The VMs are compliant.
	The VMs are not compliant.
	This is a VMware component. For example, it is an ESX/ESXi host.

Locating VMs in a Complex VM Tree

Locating VMs in the VM tree can become difficult as the VM tree grows in complexity. To simplify the process and make it easier to find specific VMs, the VM tree provides a

filter with advanced capabilities. You can enter in the filter box a text string that matches VM names within the tree. As you enter the text, the vGW dynamically searches the tree for any matches.



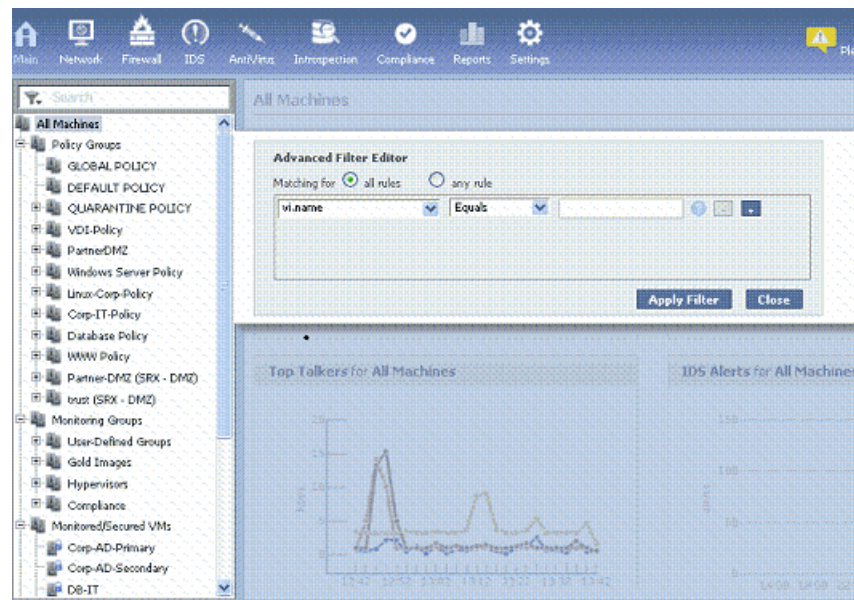
NOTE: An x icon is shown at the right side of the search field as the filter is being applied. You can use it to clear the filter.

As the filter is applied, the tree is expanded to show matching VMs. You do not need to expand all groups in the tree to find them. Branches in the tree that do not contain matches are collapsed.

You can use the Advanced Filter Editor feature to search the VM tree based on attributes rather than by name.


To use the advanced filter, click the icon at the left side of the search filter. This displays the Advanced Filter Editor shown in [Figure 35 on page 73](#).

Figure 35: Searching All VMs in the VM Tree Using the Advanced Editor



You can search based on data such as the portgroup, VLAN, and the IP protocol family using attributes such as `vi.portgroup`, `vi.vlan`, `vi.ipv4`, and `vi.ipv6`. You can also search for VMs by name. See [Figure 36 on page 74](#).

Figure 36: Searching for Specific VMs in the VM Tree Using the Advanced Editor



To remove the filter and collapse the branches, click the x icon to the right of the filter.

Related Documentation

- [Understanding the vGW Security Design VM Taskbar on page 69](#)
- [Understanding the vGW Security VM on page 64](#)
- [Understanding vGW Series on page 3](#)

PART 5

vGW Series IPv6 Support

- [IPv6 Addressing on page 77](#)
- [vGW Series IPv6 Implementation on page 85](#)
- [IPv6 and vGW Series Requirements on page 95](#)

CHAPTER 9

IPv6 Addressing

- [Understanding IPv6 Addressing on page 77](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 82](#)

Understanding IPv6 Addressing

This topic gives an overview of IP version 6 (IPv6). Then it covers the IPv6 address, including use of its header fields.

This topic includes the following sections:

- [IPv6 and the Cloud on page 77](#)
- [IPv6 and IPv4 on page 78](#)
- [IPv6 Address Space, Addressing, and Address Types on page 78](#)
- [The IPv6 Basic Packet Header on page 78](#)
- [The IPv6 Packet Header Extensions on page 80](#)
- [The IPv6 Address Format on page 81](#)
- [Address Assignment and IPv6 on page 81](#)

IPv6 and the Cloud

The ongoing expansive growth of the Internet and the need to provide IP addresses to accommodate it—including addresses for virtualized machines and resources in the cloud—is accelerating the emergent use of IPv6. IPv6 with its robust architecture was designed to support increasing numbers of new users, computer networks, Internet-enabled devices, applications for collaboration and communication, and virtualized resources. As they increase in number, applications and services within clouds render the need for transition to IPv6 even more immediate. In this and other regards, the cloud and IPv6 are intrinsic affiliates.

Whether physical or virtual, every machine requires an IP address. Because of its address size, IPv6 allows for infrastructure scalability, and the cloud allows for agility. vGW Series secures virtualized environments in the cloud and it allows for IPv6 communication. Without the scalability that IPv6 gives it, the cloud cannot extend to enable the plans and goals that are being generated for its use by companies and service providers.

As enterprise data centers and service providers undergo the transition to cloud computing, they are also evolving to support IPv6, and the two transitions are deeply related. In some cases, organizations are making the transition to the cloud and IPv6 concurrently. As they transition to the cloud, organizations and companies want to know that their data is secure. vGW Series meets these requirements in its ability to secure the virtualized network and its support of IPv6, including support for IPv4 and IPv6 dual stack, which is commonly used by companies to manage their IP transition.

IPv6 and IPv4

The number of available IPv4 addresses is limited by the IPv4 32-bit address size. IPv6, which was designed in part to fix the address limitations of IPv4, is defined by a 128-bit address size. IPv4 is widely used throughout the world today for the Internet, intranets, and private networks, but it is nearing the point where its addresses are becoming scarce and it could run out of them. IPv4 has been extended using techniques such as Network Address Translation (NAT), which allows for ranges of private addresses to be represented by a single public address, and temporary address assignment. Although useful, these techniques fall short of the requirements of environments such as virtualized networks and cloud applications, Internet-based consumer appliances, always-on systems, and continuously emerging wireless technologies.

IPv6 Address Space, Addressing, and Address Types

This section covers IPv6 addressing, and it identifies its three types of addresses. Addressing is the area where most of the differences between IPv4 and IPv6 exist, but the changes are largely about the ways in which addresses are implemented and used. IPv6 has a vastly larger address space than the impending exhausted IPv4 address space. IPv6 increases the size of the IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of the address space.

In addition to the increased address space, IPv6 differs from IPv4 in regard to addresses in the following ways. IPv6:

- Includes a scope field that identifies the type of application that the address pertains to.
- Does not support broadcast addresses, but instead uses multicast addresses to broadcast a packet.
- Defines a new type of address called anycast.

The IPv6 Basic Packet Header

This section identifies the IPv6 basic packet header fields including their bit lengths and uses. See [Table 6 on page 78](#).

Table 6: IPv6 Basic Packet Header Fields

Header Name	Bit Length	Purpose
Version	4	IPv6 version field that specifies a value of 6 indicating that IPv6 is used, as opposed to 4 for IPv4.

Table 6: IPv6 Basic Packet Header Fields (*continued*)

Header Name	Bit Length	Purpose
Traffic Class	8	Allows source nodes or routers to identify different classes (or priorities for quality of service) for IPv6 packets. (This field replaces the IPv4 Type of Service field.)
Flow Label	20	Identifies the flow to which the packet belongs. Packets in a flow share a common purpose, or belong to a common category, as interpreted by external devices such as routers or destination hosts.
Payload Length	16	Specifies the length of the IPv6 packet payload, or contents, expressed in octets.
Next Header	8	<p>Identifies the type of Internet Protocol for the header that immediately follows the IPv6 header.</p> <p>The Next Header field replaces the IPv4 Protocol field. It is an optional field. It can contain:</p> <ul style="list-style-type: none"> • an IPv6 extension header type. For example, when security is performed on exchanged packets, the Next Header value is probably 50 (ESP extension header) or 51 (AH extension header). • an upper-layer Protocol Data Unit (PDU). For example, the Next Header value could be 6 (for TCP), 17 (for UDP), or 58 (for ICMPv6). • unknown
Hop Limit	8	Specifies the maximum number of hops the packet can make.
Source IP Address	128	Identifies the host device, or interface on a host, that generated the IPv6 packet.
Destination IP Address	128	Identifies the host device, or interface on a host, to which the IPv6 packet is to be sent.

vGW Series examines the header called next-header, and if it encounters one of the following extension headers, the software parses it, and it regards the packet as belonging to the corresponding protocol:

- Internet Control Message Protocol version 6 (ICMPv6)
- Transport Control Protocol (TCP)

As part of its sanity check, vGW Series checks the TCP header length.

- UDP

As part of its sanity check, vGW Series checks the UDP header length.

- Enhanced Security Protocol (ESP) or Authentication Header (AH)



NOTE: vGW Series does not perform ESP or AH encryption.

The IPv6 Packet Header Extensions

This section defines IP version 6 (IPv6) packet header extensions.

IPv6 extension headers contain supplementary information used by network devices (such as routers, switches, and endpoint hosts) to decide how to direct or process an IPv6 packet. The length of each extension header is an integer multiple of 8 octets. This allows subsequent extension headers to use 8-octet structures.

Any header followed by an extension header contains a Next Header value that identifies the extension header type. Extension headers always follow the basic IPv6 header in order as shown in [Table 7 on page 80](#):



NOTE: The destination IP address can appear twice, once after the hop-by-hop header and another after the last extension header.

Table 7: IPv6 Extension Headers

Header Name	Purpose
Hop-by-Hop Options	Specifies delivery parameters at each hop on the path to the destination host. NOTE: A hop-by-hop option can appear only following the IPv6 basic header. If it is used, it should be the first extension header. It cannot appear after another extension header.
Destination Options	Specifies packet delivery parameters for either intermediate destination devices or the final destination host. When a packet uses this header, the Next Header value of the previous header must be 60.
Routing	Defines strict source routing and loose source routing for the packet. (With strict source routing, each intermediate destination device must be a single hop away. With loose source routing, intermediate destination devices can be one or more hops away.) When a packet uses this header, the Next Header value of the previous header must be 43.
Fragment	Specifies how to perform IPv6 fragmentation and reassembly services. When a packet uses this header, the Next Header value of the previous header must be 44. A source host uses the fragment extension header to tell the destination host the size of the packet that was fragmented so that the destination host can reassemble the packet.
Authentication	Provides authentication, data integrity, and anti-replay protection. When a packet uses this header, the Next Header value of the previous header must be 51.

Table 7: IPv6 Extension Headers (*continued*)

Header Name	Purpose
Encapsulating Security Payload	Provides data confidentiality, data authentication, and anti-replay protection for Encapsulated Security Payload (ESP) packets. When a packet uses this header, the Next Header value of the previous header must be 50.
Destination IP Address	Identifies the host device, or interface on a host, to which the IPv6 packet is to be sent. NOTE: The destination address may appear twice, the first instance after the hop limit following the source IP address and the second instance after the final extension header.

The IPv6 Address Format

This section explains the format for IPv6 addresses, including how to compress them, and it gives some examples.

All IPv6 addresses are 128 bits long, written as 8 sections of 16 bits each. They are expressed in hexadecimal representation, so the sections range from 0 to ffff. Sections are delimited by colons, and leading zeroes in each section may be omitted. If two or more consecutive sections have all zeroes, they can be collapsed to a double colon.

- IPv6 addresses have the following format in which each xxxx is a 16-bit hexadecimal value, and each x is a 4-bit hexadecimal value.

```
xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx
```

- Here is an example of an IPv6 address:

```
3ffe:0000:0000:0001:0200:f8ff:fe75:50df
```

- For an IPv6 address that contains consecutive fields of leading zeros, you can omit the zeros from each section. If you take this approach, you can write the example address that is shown previously in the following way:

```
3ffe:0:0:1:200:f8ff:fe75:50df
```

- For an IPv6 address that includes contiguous sections each of which contain zeros, vGW Series compresses the 16-bit groups of zeros to double colons (::). The double-colon delimiter can be used only once within a single IPv6 address as shown in the following example:

```
3ffe::1:200:f8ff:fe75:50df
```

Address Assignment and IPv6

The IPv6 stateless autoconfiguration feature 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.

IPv6 requires that every network interface on which the protocol is enabled have a link-local address bound to it, even when a routable address is assigned to it. Link-local

addresses are not routable. They are unique addresses in that only local traffic can be sent to them.

A link-local address is not assigned by DHCP. Consequently, IPv6 hosts often have more than one IPv6 address assigned to each of their IPv6-enabled network interfaces. Link-local addresses may be assigned statefully through mechanisms such as DHCP, but most often they are assigned using stateless autoconfiguration.

The link-local address is required for IPv6 sublayer operations of the Neighbor Discovery Protocol (NDP). NDP is an IP protocol used with IPv6 for address autoconfiguration of nodes, nodes discovery, location of routers and DNS servers, node reachability, identification of paths to active neighbor nodes, and other services related to address detection.



NOTE: You can create policies to restrict access to certain link-local addresses as required for your environment.

**Related
Documentation**

- [Understanding vGW Series IPv6 Support on page 87](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding vGW Series on page 3](#)

Understanding vGW Series IPv4 and IPv6 Dual Stack Support

This topic includes the following sections:

- [Dual Stack Background on page 82](#)

Dual Stack Background

IPv6 is designed to extend and enhance IP addressing while maintaining IPv4 functions that work well with new applications. Enterprises and service providers who convert their environments to use IPv6 often carry out the transition in phases during which some of their devices continue to use IPv4 addresses. To ensure optimum performance and a smooth transition, many companies implement a dual-stack architecture during this period. When a device has dual-stack capabilities, it has access to both IPv4 and IPv6 networks. It can use both protocols to connect to remote devices and destinations in parallel.

A dual-stack device can connect to an IPv4-only device or an IPv6-only device, or it can connect to another device that implements dual stack.



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.

For additional information on address assignment, see [“Understanding IPv6 Addressing” on page 77](#).

**Related
Documentation**

- [Understanding vGW Series IPv6 Support on page 87](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 10

vGW Series IPv6 Implementation

- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding vGW Series IPv6 Support on page 87](#)
- [IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments on page 91](#)

Overview of IPv6 Implementation in the vGW Security Design VM Modules

This topic summarizes the vGW Security Design VM IPv6 implementation that allows you to enter and view information pertaining to IPv6 addresses and traffic. It also covers information on individual modules, including figures that show windows and tabs that contain IPv6 fields and information.

This topic assumes that IPv4 addresses are handled largely in the same way.

- [Main Module on page 85](#)
- [Network Module on page 86](#)
- [Firewall Module on page 86](#)
- [IDS Module on page 87](#)
- [AntiVirus Module on page 87](#)
- [Introspection Module on page 87](#)
- [Compliance Module on page 87](#)
- [Reports Module on page 87](#)
- [Settings Module on page 87](#)

Main Module

You can view or enter IPv6 information, in addition to IPv4, in the following areas:

- Dashboard charts.
- Status. In this case, you can roll the mouse over the field to view the complete IPv6 address.
- Events and Alerts. You can use the filter box to sort the data by IPv6 addresses.
- Quarantine. You can roll the mouse over a VM name to view its IP address.

Network Module

The details tables display IPv6 information. Network traffic assessment takes into account IPv6 traffic.

Firewall Module

You can view results containing IPv6 and IPv4 addresses, and you can create policies that include them.

Firewall Logs

Firewall log entries include information pertaining to IPv6 and IPv4 addresses.

Policies

If all components belong to vGW Series release 5.5. or later, you can create firewall policies on IPv6 objects, in addition to IPv4 objects. You can select groups, networks, and machines that have IPv6 addresses to use as source and destination terms. You can also create new groups, networks, and machines that have IPv6 addresses and use them in rules.

You can use the following predefined addresses for source and destination terms in policy rules:

- Any—Any IPv4 or IPv6 address
- Any-IPv4—Any IPv4 address
- Any-IPv6—Any-IPv6 address



NOTE: Prior to vGW Series Release 5.5, which introduces support for IPv6, the predefined term “Any” referred to any IPv4 address.

ICMPv6

By default vGW Series allows a subset of Internet Control Message Protocol version 6 (ICMPv6) traffic types. These types are included in the DefaultAllow-ICMPv6 protocol group. ICMPv6 is integral to IPv6 and fundamental to the proper functioning of IPv6 networks. For more information, on vGW Series and ICMPv6 protocols, see *Understanding How vGW Series Handles ICMPv6 Protocol Traffic*.

- individual ICMPv6 protocols.
- an icmp6-all protocol definition that you can use to refer to *all* ICMPv6 protocols collectively in a policy rule.
- the DefaultAllow-ICMPv6 protocol group that includes some of the ICMPv6 protocols. DefaultAllow-ICMPv6 is used in a default inbound Global Policy rule that allows inbound traffic for the group of ICMPv6 protocols.

IDS Module

The IDS engine detects and reports attacks launched by IPv6 and IPv4 traffic.

AntiVirus Module

You use the vGW AntiVirus On-Demand and On-Access features to protect your environment from malicious attacks. The AntiVirus On-Access scan requires IPv4.

Introspection Module

You use the Introspection feature in both IPv6 and IPv4 environments. (vGW Series can mount disks that are attached to VMs that have either IPv6 or IPv4 addresses bound to them.)

Compliance Module

You can create compliance rules for hypervisors that have IPv6 or IPv4 addresses bound to them. Prebuilt compliance rules apply to both IPv6 and IPv4 environments.

Reports Module

Charts, graphics, and other areas of reports that show IPv4 addresses can also show IPv6 addresses. You can sort information based on IPv6 addresses using the filter box.

Settings Module

All sections of the Settings module that display or accept IPv4 addresses also display or accept IPv6 addresses.

Additional protocols for IPv6, including ICMPv6 protocols, an ICMPv6 transport protocol type, a protocol that includes all ICMPv6 protocols, and a default ICMPv6 protocol group that allows access for some fundamental ICMPv6 protocols have been added to the protocol list. For details, see *Understanding vGW Series Protocols Support*.

For details, see the topics that pertain to the Settings module. See *Understanding the vGW Series Settings Module*.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series IPv6 Support on page 87](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 82](#)

Understanding vGW Series IPv6 Support

This topic covers IPv6 in relation to vGW Series. It considers IPv6 with the understanding that the cloud and IPv6 are inherently linked. vGW Series secures the cloud, and it provides support for IPv6 alone or with IPv4 in a dual stack implementation.

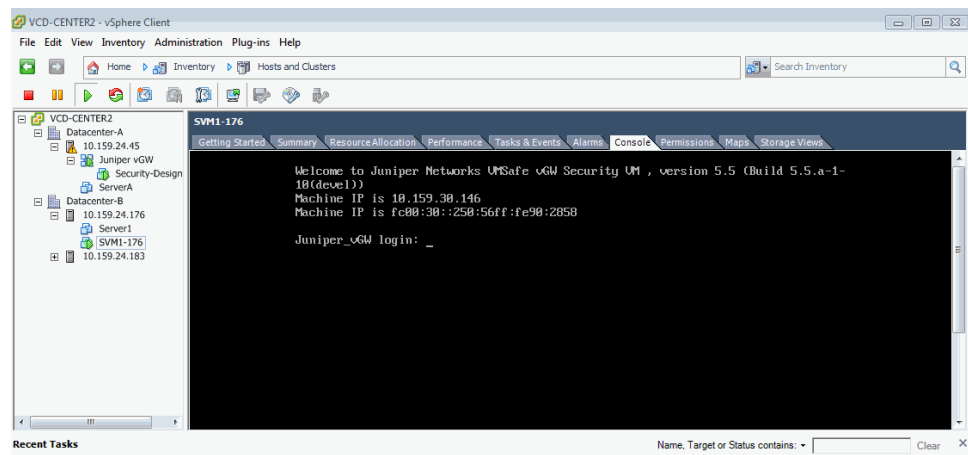
This topic covers how vGW Series displays or allows you to enter IPv6 address in the vGW Security Design VM modules.

- [vGW Security Design VM and vGW CLI Support for IPv6 Addresses on page 88](#)
- [Entering IPv6 Addresses on page 89](#)
- [vGW Series IPv6 Address Representation on page 89](#)
- [IPv4-Mapped IPv6 Addresses on page 90](#)
- [vGW Security Design VM Filter Boxes on page 90](#)
- [Searching the VM Tree for VMs and Hypervisors with IPv6 Addresses on page 90](#)
- [vGW Security Design VM and IPv6 and IPv4 Addressing on page 90](#)
- [vGW Security VM IP Addressing Support on page 90](#)

vGW Security Design VM and vGW CLI Support for IPv6 Addresses

vGW Series implements support for IPv6 addresses in all areas of the vGW Security Design VM, console, and CLI output where addresses are represented as text. Attributes used in Smart Groups that pertain to IPv4 addresses have IPv6 corollaries. For example, the `vi.ipv4` Smart Group attribute now has a `vi.ipv6` corollary. In another example, [Figure 37 on page 88](#) shows the console displaying both the IPv4 and the IPv6 addresses for the SVM-176.

Figure 37: Console Showing IPv6 and IPv4 vGW Security VM Addresses



For releases of vGW Series prior to vGW Series 5.5, IP addresses were assumed to be 32-bit IPv4 addresses. Network and host objects, security policies, and logging and reporting data were all assumed to accept or display IPv4 addresses only. For vGW Series 5.5, any area of the product that used IPv4 addresses now accepts, validates, and supports both IPv4 and IPv6 addresses. If the vGW Security Design VM is configured for dual-stack support, both IPv6 and IPv4 addresses are accepted or displayed. For details on configuring the vGW Security Design VM for dual stack, see *Configuring the vGW Series Network Settings*.



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.

Entering IPv6 Addresses

vGW Series IPv6 text representation follows the canonical text representation format for IPv6 recommended by the RFC 5952 standard. You can enter IPv6 addresses in any of the standard text representation formats, and the vGW Security Design VM will accept them as valid IPv6 addresses. However, it compresses IPv6 addresses when it displays them.

vGW Series IPv6 Address Representation

The vGW Security Design VM interface, reports, logs, log collections, CLI output, and console messages include coverage of IPv6 addresses, in addition to IPv4 addresses.

IPv6 addresses have the following format in which each `xxxx` is a 16-bit hexadecimal value, and each `x` is a 4-bit hexadecimal value.

`xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx`

Here is an example of an IPv6 address:

`3ffe:0000:0000:0001:0200:f8ff:fe75:50df`

To utilize display space, the vGW Series compresses IPv6 addresses, following the RFC 5952 standard recommendation for address compression.

For an IPv6 address that includes contiguous sections each of which contains zeros, the vGW Series compresses the 16-bit groups of zeros to double colons (`::`).

vGW Series would present the following IPv6 address that contains four sections of zeroes:

`2001:db8:0:0:0:0:2:1`

in its IPv6 compressed representation:

`2001:db8::2:1`

IPv4-Mapped IPv6 Addresses

vGW Series supports IPv4-mapped IPv6 addresses, which are a class of addresses that are utilized in hybrid dual-stack IPv6/IPv4 implementations. The first 80 bits of these addresses are zero, the next 16 bits are one, and the remaining 32 bits are the IPv4 address. In some cases, these addresses are written with the first 96 bits in standard IPv6 format and the last 32 bits written in IPv4 dot-decimal notation.

The following representation stands for the IPv4 address 192.0.2.128:

`::ffff:192.0.2.128`

vGW Security Design VM Filter Boxes

All vGW Security Design VM filters that you can use to filter on specific IP addresses and systems are enhanced to support IPv6 values.

Filter boxes for source and destination addresses display IPv4 addresses, IPv6 addresses, or both, depending on how the system is configured.

Searching the VM Tree for VMs and Hypervisors with IPv6 Addresses

You can use the VM tree search box Advanced Filter Editor to locate VMs that have IPv6 addresses bound to them. You can specify a single IPv6 address or a range of addresses to display the VMs that the addresses are assigned to. In response, vGW Series highlights the matching VMs.

vGW Security Design VM and IPv6 and IPv4 Addressing

The vGW Security Design VM supports IPv4 addresses, IPv6 addresses, and IPv4-IPv6 addresses for dual stack.

vGW Security VM IP Addressing Support

A vGW Security VM must have either an IPv4 address or an IPv6 address bound to it or both types of addresses if it is configured for dual stack. You should configure the IP address for a vGW Security VM based on how the vGW Security Design VM is configured.

Related Documentation

- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 82](#)
- [IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments on page 91](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts](#)
- [Understanding vGW Series on page 3](#)
- [Configuring vGW Series Installation Settings](#)

IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments

This topic covers how vGW Series treats the configuration of IPv6 traffic handling in homogeneous environments in which all vGW Series components—vGW Security Design VMs and vGW Security VMs—belong to vGW Series 5.5 (or later) and heterogeneous environments in which they do not. A heterogeneous environment might include a 5.5 vGW Security Design VM that manages one or more 5.0 vGW Security VMs.

This topic includes the following sections:

- [IPv6 Traffic Handling in Homogenous Environments \(All vGW Series Components at Version 5.5 or Later\)](#) on page 91
- [IPv6 Traffic Handling in Heterogeneous Environments \(with a Mix of vGW Series Component Versions\)](#) on page 91

IPv6 Traffic Handling in Homogenous Environments (All vGW Series Components at Version 5.5 or Later)

If your environment contains a mix of vGW Series components with different versions because it is in a transition period, skip this section and read [“IPv6 Traffic Handling in Heterogeneous Environments \(with a Mix of vGW Series Component Versions\)”](#) on page 91.

If your environment is a complete vGW Series 5.5 installation, you can create granular firewall policies on IPv6 traffic flows. For example, you can create policies that use IPv6 objects such as IPv6 machines or the predefined term Any-IPv6, which pertains exclusively to IPv6 traffic. To do so, you use the Firewall module Manage Policy page, just as you would do to create rules for a given policy for IPv4 traffic.

For a complete vGW Series 5.5 installation, for the IPv6 traffic configuration option, the Allow check box in the Global settings is dimmed, and it is not used.



NOTE: A complete installation of vGW Series 5.5 on all components is required to take advantage of the ability to write granular policies on IPv6 traffic flows.

IPv6 Traffic Handling in Heterogeneous Environments (with a Mix of vGW Series Component Versions)

vGW Series enables support of IPv6 in environments that include a mix of vGW Series 5.5 or later components and vGW Series 5.0 components. This kind of environment is not uncommon during the transition period when organizations are adopting IPv6 but continue to use IPv4 until the transition is completed.

A heterogeneous environment might include any combination of vGW components with different versions. For example, an environment might include:

- 5.5 vGW Security Design VMs and one or more 5.0 vGW Security VMs.

- 5.5 vGW Security VMs, a 5.5 vGW Security Design VM, and a 5.0 vGW Security Design VM.



NOTE: Until all components in your environment are at version 5.5 or later, you must use the Settings module Security Settings > Global page **IPv6 traffic** configuration option to control handling of IPv6 traffic.

You can continue to create granular IPv4 policies and push them to vGW Security VMs.

After you upgrade all vGW Security VMs in your environment to vGW Series 5.5, the Security Settings > Global page **IPv6 traffic** setting is no longer used. Instead, policy rules are applied. In this case, the behavior might be different from what you expect if you presume that the global setting is still in effect.

It is important to understand how vGW Series treats heterogeneous environments in regard to IPv6. For vGW Series 5.5 or later environments in which not all components have been upgraded to version 5.5:

- Traffic mirroring and IDS for IPv6 is not available.
- You cannot control IPv6 traffic at the granular policy rule level. For example, you cannot create firewall policies that use IPv6 objects, such as IPv6 machines or the predefined term **Any-IPv6** that pertains exclusively to IPv6 traffic. If you attempt to do so, vGW Series does not allow you to save the policy. vGW Series displays the following error message when you attempt to save the policy:

“There are vGW components older than version 5.5. Granular IPv6 policy creation is not supported when any vGW component is older than version 5.5. IPv6 traffic will be either accepted or dropped based on Settings > Security > Global > IP Traffic configuration option until all components are upgraded.”

- The predefined term **Any** matches on only IPv4 traffic, as it does in releases prior to vGW Series 5.5.
- vGW Series 5.5 relies on the configuration of the **IPv6 traffic** Global setting to determine whether to allow or drop IPv6 traffic. In this case, the behavior is unchanged from that of vGW Series 5.0 and earlier releases.

Only the **IPv6 traffic** setting is used, and it is global in that it applies to all IPv6 traffic.



NOTE: When you upgrade a vGW Security Design VM to version 5.5 from a preceding release, vGW Series carries over and continues to use the **IPv6 traffic** setting from the previous release.

Related Documentation

- [Understanding vGW Series IPv6 Support on page 87](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 82](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)

- [Understanding vGW Series on page 3](#)

CHAPTER 11

IPv6 and vGW Series Requirements

- [Understanding Requirements for Communication Between vGW Series IPv6 Components and Juniper Networks IPv4 Servers on page 95](#)
- [Understanding vGW Series IPv6 Accessibility Requirements for Customer Servers on page 96](#)

Understanding Requirements for Communication Between vGW Series IPv6 Components and Juniper Networks IPv4 Servers

When a vGW Series component with an IPv6 address must communicate with a Juniper Networks server that is on an IPv4 network, an incompatibility exists. To overcome the IP protocol incompatibility, the following capabilities and implementations must exist:

- Connectivity to the IPv4 Internet.
- An address translation mechanism such as NAT64.
- A DNS solution such as DNS64 that provides an authentication, authorization, accounting, and address (AAAA) record with an IPv6 address through which the server can be reached.

The Juniper Networks servers that vGW Series communicates with that could be affected by this compatibility issue are:

- [updates.altornetworks.com](#)—Update server.
- [anubis.altornetworks.com](#)—Log collection upload server.
- [update.juniper-updates.net](#)—vGW AntiVirus signature updates server. This server is accessed only by the vGW Security Design VM.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Configuring Global Settings Using the vGW Series Settings Module \(VMware\)](#)
- [Understanding vGW Series IPv6 Support on page 87](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding vGW Series Log Collection](#)

Understanding vGW Series IPv6 Accessibility Requirements for Customer Servers

When vGW Series components with IPv6 addresses bound to them must communicate with customer servers that are accessible only through IPv4, compatibility issues might arise that require your consideration. These servers include:

- Syslog
- Netflow
- Mirroring (GRE)
- AD server
- SMTP server
- SRX - for zones integration
- DNS server

Related Documentation

- [Understanding vGW Series IPv6 Support on page 87](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 85](#)
- [Understanding Requirements for Communication Between vGW Series IPv6 Components and Juniper Networks IPv4 Servers on page 95](#)
- [Understanding vGW Series on page 3](#)
- [*Adding and Editing vGW Series Machines Definitions \(VMware\)*](#)

PART 6

Index

- [Index on page 99](#)

Index

Symbols

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

B

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

C

cloud computing.....	6
comments, in configuration statements.....	xiii
conventions	
text and syntax.....	xii
curly braces, in configuration statements.....	xiii
customer support.....	xiv
contacting JTAC.....	xiv

D

DNS services.....	31
documentation	
comments on.....	xiii
dual stack.....	82

E

ESX/ESXi hosts.....	7
events and alerts.....	16

F

font conventions.....	xii
-----------------------	-----

H

hypervisor	
vGW Series kernel module.....	8
hypervisors	
overview.....	8

I

installation	
prerequisites and resource requirements.....	31
IPv6	
address.....	77
communication with Juniper Networks IPv4	
servers.....	95
dual stack, IPv6 and IPv4.....	82
header fields.....	77
mixed component versions.....	91
overview.....	77
vGW Security Design VM modules.....	85

K

kernel module.....	8
--------------------	---

M

Main module.....	16
manuals	
comments on.....	xiii

N

network connectivity.....	31
NTP.....	36
NTP services.....	31

O

Open Virtualization Format (OVF)	
OVA template.....	39, 40
OVA bundled method.....	40
downloading package.....	41
OVA template.....	39
OVA template method	
single method for vGW Security Design	
VM.....	49
single OVA install method for vGW Security	
VM.....	51
overview.....	3

P

parentheses, in syntax descriptions.....	xiii
port groups.....	31

R

resource requirements	
DNS services.....	31
network connectivity.....	31
NTP services.....	31
port groups.....	31

vCenter.....	31	integrating vGW Series	
virtual appliances.....	31	preparation.....	35
virtual devices.....	31	VMotion.....	7
vNICs.....	31	vSphere.....	7
vSphere.....	31	vSphere operating system.....	7, 31
vSwitches.....	31	<i>See also</i> Virtual Center (vCenter)	
Web browsers.....	31	VMware VMsafe APIs.....	36
		vNICs.....	31
S		vSphere.....	7
status and status icons.....	16	vSwitches.....	31
support, technical <i>See</i> technical support			
syntax conventions.....	xii	W	
		Web browsers supported.....	31
T			
taskbar.....	69		
technical support			
contacting JTAC.....	xiv		
time synchronization.....	36		
V			
vCenter.....	31		
vGW Security Design			
modules.....	11		
vGW Security Design VM			
in mixed component version IPv6			
environment.....	91		
IPv6.....	85		
Main module.....	16		
navigation.....	67		
overview.....	63		
summary.....	3		
taskbar.....	69		
vGW Security VM			
in mixed component version IPv6			
environment.....	91		
installation modes.....	36		
overview.....	64		
single OVA install method.....	51		
summary.....	3		
virtual appliances.....	31		
virtual devices			
sizes.....	31		
VMotion.....	7		
VMsafe Firewall + Monitoring mode.....	36		
VMsafe Monitoring mode.....	36		
VMware.....	31		
ESX/ESXi hosts.....	7		
infrastructure and vGW.....	7		