



vGW Series

Installation and Administration Guide for VMware

Release

5.5r4



Published: 2013-11-04

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

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 Installation and Administration Guide for VMware
Release 5.5r4
Copyright © 2013, Juniper Networks, Inc.
All rights reserved.

Revision History
July 2013—vGW Series 5.5r4

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

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.

Abbreviated Table of Contents

	About This Guide	xxiii
Part 1	vGW Series Overview	
Chapter 1	Introduction to vGW Series	3
Chapter 2	vGW Series Resource Requirements	15
Chapter 3	vGW Series Integration with the VMware Environment	21
Chapter 4	vGW Series IPv6 Support	49
Part 2	vGW Series Configuration and Management	
Chapter 5	The vGW Security Design VM and the vGW Security VM	67
Chapter 6	vGW Series Main Module	77
Chapter 7	vGW Series Network Module	85
Chapter 8	vGW Series Firewall Module	93
Chapter 9	vGW Series IDS Module	121
Chapter 10	vGW Series AntiVirus Module	131
Chapter 11	vGW Series Introspection Module	163
Chapter 12	vGW Series Compliance Module	183
Chapter 13	vGW Series Reports Module	193
Chapter 14	vGW Series Settings Module	203
Chapter 15	vGW Series Application Settings	205
Chapter 16	vGW Series Security Settings	281
Chapter 17	vGW Security Design Appliance Settings	319
Chapter 18	vGW Series Status Alerts	339
Chapter 19	High Availability and Fault Tolerance	343
Part 3	Juniper Networks Products Interoperability	
Chapter 20	vGW Series Interoperability with Juniper Networks Products	359
Part 4	Index	
	Index	371

Table of Contents

	About This Guide	xxiii
	vGW Series Documentation	xxiii
	Objectives	xxiii
	Audience	xxiii
	Documentation Conventions	xxiv
	Obtaining Documentation	xxiv
	Documentation Feedback	xxiv
	Requesting Technical Support	xxiv
	Self-Help Online Tools and Resources	xxv
	Opening a Case with JTAC	xxv
Part 1	vGW Series Overview	
Chapter 1	Introduction to vGW Series	3
	Understanding vGW Series	3
	Understanding the vGW Series Architecture	4
	vGW Security Design VM Modules (VMware)	6
	Understanding Cloud Computing and vGW Series	11
	Understanding Hypervisors and vGW Series	12
	Understanding the VMware Infrastructure and vGW Series	13
	Understanding vSphere and the vGW Series	13
	Understanding VMware ESX and ESXi Hosts and the vGW Series	13
	Understanding VMotion and vGW Series	14
Chapter 2	vGW Series Resource Requirements	15
	vGW Series Prerequisites and Resource Requirements for the VMware	
	Environment	15
	Overall Resource and Access Requirements	15
	Virtual Appliance System Requirements	16
	vGW Series VMware vSwitch Requirements	17
	VMware Port Group Requirements	18
	Virtualized NIC Requirements	18
Chapter 3	vGW Series Integration with the VMware Environment	21
	Preparing to Integrate the vGW Series with the VMware Environment	21
	vGW Series VMsafe Firewall + Monitoring and VMsafe Monitoring Modes	22
	Understanding vGW Series Environment Time Synchronization	23

	Understanding the Open Virtualization Format OVA Template Method	23
	Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure	24
	Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware	34
	Using the OVA Single File Method to Integrate the vGW Security VM with VMware	36
	Setting Up vGW Series	36
	Determining the vGW Security Design VM's Default IP Address	37
	Changing or Setting the IP Address for the vGW Security Design VM	38
	Connecting to the vGW Security Design VM and Configuring Basic Settings	39
Chapter 4	vGW Series IPv6 Support	49
	Understanding IPv6 Addressing	49
	IPv6 and the Cloud	49
	IPv6 and IPv4	50
	IPv6 Address Space, Addressing, and Address Types	50
	The IPv6 Basic Packet Header	51
	The IPv6 Packet Header Extensions	52
	The IPv6 Address Format	53
	Address Assignment and IPv6	54
	Understanding vGW Series IPv6 Support	54
	vGW Security Design VM and vGW CLI Support for IPv6 Addresses	55
	Entering IPv6 Addresses	56
	vGW Series IPv6 Address Representation	56
	IPv4-Mapped IPv6 Addresses	57
	vGW Security Design VM Filter Boxes	57
	Searching the VM Tree for VMs and Hypervisors with IPv6 Addresses	57
	vGW Security Design VM and IPv6 and IPv4 Addressing	57
	vGW Security VM IP Addressing Support	57
	Overview of IPv6 Implementation in the vGW Security Design VM Modules	58
	Main Module	58
	Network Module	58
	Firewall Module	58
	Firewall Logs	58
	Policies	58
	ICMPv6	59
	IDS Module	59
	AntiVirus Module	59
	Introspection Module	59
	Compliance Module	59
	Reports Module	60
	Settings Module	60

	IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments	60
	IPv6 Traffic Handling in Homogenous Environments (All vGW Series Components at Version 5.5 or Later)	60
	IPv6 Traffic Handling in Heterogeneous Environments (with a Mix of vGW Series Component Versions)	61
	Understanding vGW Series IPv4 and IPv6 Dual Stack Support	62
	Dual Stack Background	62
	Understanding Requirements for Communication Between vGW Series IPv6 Components and Juniper Networks IPv4 Servers	63
	Understanding vGW Series IPv6 Accessibility Requirements for Customer Servers	64
Part 2	vGW Series Configuration and Management	
Chapter 5	The vGW Security Design VM and the vGW Security VM	67
	Understanding the vGW Security Design VM	67
	Understanding the vGW Security VM	68
	Understanding vGW Security Design VM Navigation	69
	Understanding the vGW Series Kernel Module	71
	Understanding the vGW Security Design VM Taskbar	71
	About the vGW Security Design VM Tree	73
	VM Tree Overview	73
	Locating VMs in a Complex VM Tree	75
Chapter 6	vGW Series Main Module	77
	Understanding the vGW Series Main Module	77
	Dashboard	77
	Status Tab	78
	Events and Alerts Tab	80
	Security Alerts	81
	System Status and Events	82
	Quarantine Tab	83
Chapter 7	vGW Series Network Module	85
	Understanding the vGW Series Network Module	85
	Network Module	85
	Manipulating Displayed Information	86
	Changing the Time Interval for Displayed Information	87
	Using Advanced Options for Filtering Network Data	89
	Sorting Table Data	89
	Using the vGW Series Network and Firewall Modules Cooperatively	90
	Network Assessment	90
	Using the Network Module to Observe Traffic Coming Into and Going Out from VMs	91
	Detecting Unexpected and Unwanted Behavior	91
	Using the Network and Firewall Modules Together	92

Chapter 8	vGW Series Firewall Module	93
	Understanding the vGW Series Firewall Module	93
	The Firewall Module and the VM Tree	93
	Overview of the Firewall Policy Model	94
	Global Policy, Group Policy, and Individual VM Policy Tiers	95
	Global Policy	95
	Group Policy	97
	Individual VM Policy Rules	97
	Default Policy	98
	Quarantine Policy	98
	Firewall Policy Structure and Policy Rules Precedence	98
	Viewing the Complete Policy Rule Base for a VM	100
	The Manage Policy Tab	100
	Policy Per vNIC and Dual Stack	101
	Creating a Policy Rule	101
	The Apply Policy Tab	104
	The Logs Tab	105
	Understanding How vGW Series Handles ICMPv6 Protocol Traffic	106
	About ICMPv6	107
	Filtering ICMPv6 Packets	107
	Default Policy Group for Allowing Inbound ICMPv6 Packets	108
	Viewing the Default ICMPv6 Protocols Group Members	108
	Editing the Default ICMPv6 Protocols Group Members	109
	Understanding Predefined Objects for vGW Series Firewall Policy Terms	110
	Defining and Selecting Source and Destination Terms for Policy Rules	110
	Predefined Global IP Address Objects	111
	Predefined Network Objects	111
	Predefined Network Objects for Well Known IP Addresses	112
	Additional IPv4 and IPv6 Predefined Network Objects	112
	Configuring vGW Series Firewall Policies	114
	Understanding vGW Series Predefined Firewall Policy for Its Components	120
Chapter 9	vGW Series IDS Module	121
	Understanding the vGW Series IDS Module	121
	Managing and Sorting Displayed Alerts Information	121
	Top Alerts Page	122
	Alert Sources Page	127
	Alert Targets Page	127
	All Alerts Page	127
	Configuring IDS Settings and Viewing Activity	128
Chapter 10	vGW Series AntiVirus Module	131
	Understanding vGW Series AntiVirus	131
	About Antivirus Software	132
	Signature-Based Detection	132

	The vGW AntiVirus Feature	133
	The vGW AntiVirus Dashboard	135
	vGW AntiVirus Configuration Overview	139
	Configuring vGW Series AntiVirus On-Access Scanning	145
	Understanding Quarantined VMs and Files Resulting from a vGW AntiVirus On-Access Scan	148
	Understanding and Installing the vGW Endpoint	149
	Installing the vGW Endpoint	149
	vGW AntiVirus Endpoint Auto-Update	150
	vGW Endpoint on the VM	150
	Quarantined Files	152
	vGW Endpoint Components and Displays	152
	vGW Endpoint Behavior	153
	Configuring vGW Series AntiVirus On-Demand Scanning	153
	Understanding Quarantined VMs and How to Manage Them	157
	About vGW Series Quarantine	158
	Configuring a Quarantine Policy	159
	Viewing the Quarantined VMs, Releasing Them From Quarantine, and Resolving Problems	160
Chapter 11	vGW Series Introspection Module	163
	Understanding the vGW Series Introspection Module	163
	Understanding the vGW Series Introspection Applications Tab	165
	Understanding the vGW Series Introspection VMs Tab	168
	Understanding the vGW Series Introspection Image Enforcer Feature	169
	Understanding the vGW Series Image Enforcer Tab	171
	Understanding the vGW Series Enforcer Profiles Tab	172
	About the Enforcer Profiles Screen	172
	The Add Enforcer Profile Pane	173
	Understanding the vGW Series Introspection Scheduling Feature	176
	Understanding the vGW Series Introspection Scan Status	177
	Understanding the vGW Series Introspection Registry Check Feature	178
	Configuring the vGW Series Introspection Registry Feature	180
Chapter 12	vGW Series Compliance Module	183
	Understanding the vGW Series Compliance Module	183
	The Compliance Module	183
	The Compliance Tab	184
	The Rules Tab	185
	Configuring a Compliance Rule	185
	Understanding the vGW Series Hypervisor and Extended VM Security	188
	The Need for Hypervisor Security	189
	vGW Series Hypervisor and VM Security, and VMware Hardening Guidelines	189
	vGW Series Hypervisor and VM Security Overview	189
	Remediation	190
	Configuration Example	190

Chapter 13	vGW Series Reports Module	193
	Understanding the vGW Series Reports Module	193
	Configuring a vGW Series Report	195
	Configuring Specifications for Automated Reports Using the vGW Series Reports Module	198
	Understanding vGW Series Custom Report Types	199
	Understanding vGW Series Network Reports	199
	About the vGW Series Firewall Reports	200
	About the vGW Series IDS Reports	200
	About the vGW Series Introspection Reports	201
	Understanding the vGW Series Compliance Report	201
	Understanding the vGW AntiVirus Report	202
Chapter 14	vGW Series Settings Module	203
	Understanding the vGW Series Settings Module	203
Chapter 15	vGW Series Application Settings	205
	Understanding the vGW Series Application Settings	206
	Viewing Status and License Information Using the vGW Series Settings Module	207
	Understanding Licenses for the vGW Series	208
	License Requirements	208
	vGW Series Licenses	208
	Evaluation Licenses	209
	Obtaining, Installing, and Managing vGW Series Licenses	209
	Integrating the vGW Series with VMware Using the Settings Module	209
	Understanding vGW Series Integration with vCloud Director	214
	VMware vCloud Director	214
	vGW Series and vCloud	214
	Requirements	215
	Configuring vGW Series Integration with vCloud Director	215
	Installing vGW Security VMs on ESX/ESXi Hosts	218
	Configuring vGW Series Installation Settings	224
	Understanding VMware Auto Deploy for ESXi Servers and with vGW Series	225
	About VMware Auto Deploy	225
	vGW Series Support for Auto Deploy	226
	vGW Series Automatic Installation of vGW Security VMs	226
	Configuring VMware Auto Deploy and vGW Series to Secure ESXi Hosts	227
	Configuring Auto Deploy in VMware	227
	Configuring vGW Series Auto Deploy Support	232
	Understanding vGW Series Timeout Parameters and the vGW Security VM	
	Installation, Uninstallation, and Update Tasks	234
	Removing vGW Security VMs from ESX/ESXi Hosts	235
	Securing and Unsecuring Virtual Machines Using the vGW Security Design VM	237
	Disabling the vGW Series Suspend-Resume Process Enacted After a VM Is Unsecured	237
	Displaying the State of the vmsafe config Setting	237
	Disabling the Suspend-Resume Process	238

Understanding Automatic Securing of VMs	238
Understanding the vGW Series Split-Center Feature	240
Understanding the Multi-Center Feature	244
The Multi-Center Feature	245
Deploying vGW Series in an Environment With a Mix of Delegate and Stand-alone vGW Security Design VMs in Various vCenters	246
Configuring vGW Series Multi-Center	247
vGW Security Design VM Master Center	247
vGW Security Design VM Delegate Centers	247
Configuring Multi-Center	248
Understanding vGW Series Multi-Center Synchronized Objects	250
Object Synchronization	250
Object Naming	251
Creation of Objects Local to the Delegate vGW Security VM	251
Configuring Scaling Using the Multi-Center and Split-Center Features	251
Understanding the vGW Series Policy per vNIC Feature	258
About Policy per vNIC	258
Why Use Policy per vNIC	259
vNICs With Individual Polices and Smart Groups	259
Viewing vNICs With Individual Policies	259
Naming Conventions for vNICs	260
Configuring the vGW Series Policy per vNIC Feature	261
Configuring Policy per vNIC to Secure Only Some of a VM's vNICs	263
Configuring and Displaying vGW Policies for Individual vNICs on the Same VM	263
Understanding Policy per vNIC and Smart Groups for VMware Environments	266
Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module	269
Configuring an Administrator Account	269
Changing Administrator Passwords	272
Global Administrator: Changing Your Own Password	272
Global Administrator: Changing the Password of Another Administrator	273
VM Administrator and Network Monitoring Administrator Accounts: Changing Your Own Password	274
Setting Up Active Directory for vGW Series Administrator Authentication	275
Adding and Editing vGW Series Machines Definitions (VMware)	277
Adding a Machine	277
Viewing Machine Information	279
Configuring vGW Series E-Mail and Reporting Applications Settings	280
Chapter 16 vGW Series Security Settings	281
Understanding the vGW Series Security Settings	281
Configuring Global Settings Using the vGW Series Settings Module (VMware)	282
vGW Series Global Settings Overview	282
Global Settings	283
vGW Series IPv6 Support and Global Settings	285
Understanding the vGW Security VM Settings	286

	Understanding and Configuring the vGW Series AntiVirus Settings	291
	Understanding and Configuring IDS Settings	292
	Understanding and Configuring IDS Signatures Settings	295
	Understanding the vGW Series Security Alert Settings	297
	Event Types	297
	E-mail Alert Settings	298
	SNMP Trap Settings	298
	AutoConfig and Multicast Alerts	298
	Understanding vGW Series Protocols Support	299
	The Protocols Page and Table	299
	Creating Protocol Groups	299
	ICMPv6	299
	Additional Protocols Added for IPv6	300
	Understanding vGW Series Groups	300
	Uses of Groups	301
	vGW Series Group Types	301
	Policy Groups and Monitoring Groups	301
	Defining the Group as a Policy Group Option with Automatic or Manual Selected	302
	Copying Groups	302
	Automatically Applying Policy Rules to VMs in Policy Groups	303
	Understanding vGW Series Smart Groups	305
	Background	305
	About Smart Groups	306
	About Using vGW Series Attributes for VMware	306
	Creating vGW Series Smart Groups for VMware	307
	vGW Series Attributes for VMware	311
	Understanding the Settings Module Networks Settings	316
	Understanding the vGW Series SRX Zones Settings	317
Chapter 17	vGW Security Design Appliance Settings	319
	Updating the vGW Security Design VM	319
	Updating Individual vGW Security VMs	321
	Updating vGW Security VMs in Batch Mode	322
	Configuring the vGW Series Network Settings	323
	The Network Configuration Page	324
	Changing the Host Name and DNS Settings	324
	Configuring Addresses for the vGW Security Design VM Interface for Communication With vGW Security VMs	325
	Changing the Way vGW Security Design VM Acquires Its Interface 1 IP Addresses	325
	Configuring the vGW Security Design VM Not to Use Dual Stack	326
	Configuring vGW Series Proxy Settings	328
	Configuring vGW Series Time Settings	328
	Understanding the vGW Series Backup and Restore Feature	329
	Configuring the vGW Series Backup and Restore Feature	331
	Understanding vGW Series Log Collection	333
	Log Collection	334
	Generating the Log Collections	334

	Uploading the File	335
	Downloading the File	336
	Using a Method Other Than the vGW Security Design VM to Generate Log Collections for It	336
	Viewing the vGW Series Logs	336
	Understanding vGW Series Support Settings	336
Chapter 18	vGW Series Status Alerts	339
	Understanding vGW Series Status and Alerts	339
	Status	339
	Alerts	339
	E-Mail Alert Settings	340
	SNMP Trap Settings	340
	AutoConfig and Multicast Alerts	340
Chapter 19	High Availability and Fault Tolerance	343
	Understanding the vGW Series High Availability Solution	343
	vGW Series HA	343
	vGW Series HA and VMware HA	344
	vGW Series HA for the vGW Security Design VM	344
	vGW Security Design VM HA Behavior	346
	vGW Series HA for the vGW Security VM	348
	Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability	349
	Installing a Secondary vGW Security VM for High Availability	353
	Understanding vGW Series Fault Tolerance Support	354
	About vGW Series Fault-Tolerance	354
	vGW Series Fault Tolerance in the vGW Series	355
	Enabling Fault Tolerance for a Virtual Machine	355
Part 3	Juniper Networks Products Interoperability	
Chapter 20	vGW Series Interoperability with Juniper Networks Products	359
	vGW Series and SRX Series Security Zones	359
	About SRX Series Services Gateway Security Zones	359
	SRX Series Services Gateway Zones and the vGW Series	360
	Enabling the Junoscript Interface for vGW Series	360
	Configuring Zone Objects for vGW Series Interoperability with SRX Series Devices	362
	About Populating vGW Series Records to SRX Series Zone Address Books	363
	Validating vGW Series Interoperability with SRX Series Zones	363
	Configuring vGW Series to Send Syslog and Netflow Data to Juniper Networks STRM Series Devices	364
	Configuring the vGW Series and IDP Series Inter-Operation	367
Part 4	Index	
	Index	371

List of Figures

Part 1	vGW Series Overview	
Chapter 1	Introduction to vGW Series	3
	Figure 1: vGW Series Architecture	5
	Figure 2: Main Module	7
	Figure 3: Network Module	7
	Figure 4: Firewall Module	8
	Figure 5: IDS Module	8
	Figure 6: AntiVirus Module	9
	Figure 7: Introspection Module	9
	Figure 8: Compliance Module	10
	Figure 9: Reports Module	10
	Figure 10: Settings Module	11
Chapter 3	vGW Series Integration with the VMware Environment	21
	Figure 11: OVA Template Details Page	26
	Figure 12: OVA File Deployment License Agreement	27
	Figure 13: Naming the vApp	28
	Figure 14: Specifying the Host and Cluster	29
	Figure 15: Selecting the Storage	29
	Figure 16: Mapping the vGW Management Networks	30
	Figure 17: Specifying the Database Disk Size	31
	Figure 18: Verifying That the Configuration Is Correct	31
	Figure 19: Displaying the vGW Appliance Components	32
	Figure 20: vGW Security Design VM Summary Tab in vCenter	34
	Figure 21: Viewing the vGW Security Design VM IP Address in VMware	38
	Figure 22: vGW Security Design VM IP Addresses on the vGW Series CLI Console	38
	Figure 23: Configuring an IP Address for the vGW Security Design VM	39
	Figure 24: Logging In to the vGW Security Design VM	40
	Figure 25: vGW Series Installation Wizard Overview	40
	Figure 26: Changing the Default Password	41
	Figure 27: Configuring Network Settings for the vGW Security Design VM	43
	Figure 28: Configuring the Time Server	44
	Figure 29: Configuring vGW Series Product Licensing	44
	Figure 30: Configuring Permanent Licenses	45
	Figure 31: Configuring the vGW Security Design VM vCenter Settings	46
Chapter 4	vGW Series IPv6 Support	49
	Figure 32: Console Showing IPv6 and IPv4 vGW Security VM Addresses	55

Part 2	vGW Series Configuration and Management	
Chapter 5	The vGW Security Design VM and the vGW Security VM	67
	Figure 33: Main Module Displayed at Login	68
	Figure 34: vGW Security Design VM Taskbar	69
	Figure 35: VM Tree	70
	Figure 36: vGW Security Design VM Taskbar	71
	Figure 37: VM Tree with Selected VMs	73
	Figure 38: Searching All VMs in the VM Tree Using the Advanced Editor	76
	Figure 39: Searching for Specific VMs in the VM Tree Using the Advanced Editor	76
Chapter 6	vGW Series Main Module	77
	Figure 40: Dashboard Tab	78
	Figure 41: Status Tab	79
	Figure 42: Taskbar Showing the Health Status Icon	80
	Figure 43: Main Module Events and Alerts Page	81
	Figure 44: Consolidated Logs for Events and Alerts	81
	Figure 45: Quarantine Tab	83
Chapter 7	vGW Series Network Module	85
	Figure 46: Network Summary Tab for All VMs	86
	Figure 47: Main Module Network Module Summary Tab for a Single VM	86
	Figure 48: Displaying Network Data for Different Time Intervals: Part 1	87
	Figure 49: Displaying Network Data for Different Time Intervals: Part 2	87
	Figure 50: Selecting a Time Interval	88
	Figure 51: Setting the Custom Time Period	88
	Figure 52: Top Protocols Across All Machines Example	90
	Figure 53: Network Module Connection Tab Information	91
Chapter 8	vGW Series Firewall Module	93
	Figure 54: Firewall Module Policy for a Single VM	94
	Figure 55: Global Policy	97
	Figure 56: VM Policy Expanded Rule Base	100
	Figure 57: Firewall Module Manage Policy Page	101
	Figure 58: Adding a Rule	102
	Figure 59: Using the Dialog Box Filter to Add Terms for policy rules	102
	Figure 60: Firewall Apply Policy Page	104
	Figure 61: Firewall Module Logs Tab	106
	Figure 62: Default Global Policy Showing Default ICMPv6 Allow Group	108
	Figure 63: Protocols Settings ICMPv6 Default Protocol Group	108
	Figure 64: Default Global Policy	115
	Figure 65: Adding a Global Policy Rule to Reject Telnet Connection Attempts	116
	Figure 66: VM Policy for an Individual VM	118
	Figure 67: Complete VM Policy for an Individual VM	118
Chapter 9	vGW Series IDS Module	121
	Figure 68: IDS Top Alerts	122
	Figure 69: IDS Top Alerts Advanced Options	123
	Figure 70: IDS Alert Description	124
	Figure 71: IDS Alert Details	124

	Figure 72: IDS Alert Details Showing Affected Systems	125
	Figure 73: IDS Alert Sources	126
	Figure 74: IDS Alert Targets	126
	Figure 75: IDS All Alerts	127
	Figure 76: IDS Settings Page	128
Chapter 10	vGW Series AntiVirus Module	131
	Figure 77: vGW AntiVirus Dashboard	135
	Figure 78: Virus Alerts	137
	Figure 79: vGW AntiVirus Scanner Config Tab	137
	Figure 80: Quarantined Files	138
	Figure 81: On-Access Scan	140
	Figure 82: vGW AntiVirus Dashboard	142
	Figure 83: Scanner Config Tab	143
	Figure 84: AntiVirus On-Access Quarantined Files	148
	Figure 85: vGW AntiVirus Settings	150
	Figure 86: vGW AntiVirus Endpoint Connection Process Dialog Box	151
	Figure 87: vGW AntiVirus Endpoint Threat Detection Dialog Box	151
	Figure 88: Scanner Config Tab	155
	Figure 89: Step 2: Scan Schedule	155
	Figure 90: Quarantine Policy in the VM Tree	159
	Figure 91: Configuring a vGW Series Quarantine Policy	160
	Figure 92: Main Module Quarantine Tab	160
Chapter 11	vGW Series Introspection Module	163
	Figure 93: vGW Series Introspection Module Applications Tab	165
	Figure 94: vGW Series Introspection Module VMs Tab	168
	Figure 95: vGW Series Introspection Module Image Enforcer Tab	171
	Figure 96: vGW Series Introspection Module Enforcer Profiles Tab	173
	Figure 97: Adding a vGW Series Introspection Module Image Enforcer Profile	173
	Figure 98: Introspection Module Scheduling Page	176
	Figure 99: vGW Series Introspection Module Scan Status Page	178
	Figure 100: Configuring a New Registry Key	180
	Figure 101: Add Schedule for Scan Page	181
	Figure 102: Add an Enforcer Profile that Allows for Registry Scans	182
Chapter 12	vGW Series Compliance Module	183
	Figure 103: vGW Series Compliance Module	184
	Figure 104: vGW Series Compliance Module Rules Tab	185
	Figure 105: Adding a Predefined Compliance Rule	188
Chapter 13	vGW Series Reports Module	193
	Figure 106: Adding a vGW Series Report Using the Reports Module	196
	Figure 107: Defining General, Destination, and Scheduling Information for the Report	196
	Figure 108: Configuring the Report Destination and Generation Schedule	197
	Figure 109: Configuring the Types of Reports to Generate	197
Chapter 14	vGW Series Settings Module	203
	Figure 110: vGW Series Settings Module	203
Chapter 15	vGW Series Application Settings	205

	Figure 111: VGW Security Design VM vCenter Integration Window Showing vCloud Director Settings Pane	217
	Figure 112: Securing an ESX/ESXi Host With a vGW Security VM	219
	Figure 113: Installing a vGW Security VM on an ESX/ESXi Host	220
	Figure 114: Specifying vGW Security Parameters During Installation	220
	Figure 115: vGW Security VM Installation Process Completion Notice	223
	Figure 116: vGW Series Failure Alert	226
	Figure 117: Configuring Automatic Installation of vGW Security VMs for Auto-Deployed ESXi Hosts	232
	Figure 118: vGW Series CLI Console	235
	Figure 119: vGW Security VM Uninstall	236
	Figure 120: Configuring the Management Scope During Installation to Include Clusters	242
	Figure 121: vGW Series Multi-Center	246
	Figure 122: Multi-Center Configuration Page at Master vGW Security Design VM	248
	Figure 123: Delegate Center Configuration on the Master vGW Security Design VM	249
	Figure 124: Delegate Center Configuration on the Master vGW Security Design VM	254
	Figure 125: Policy for Single vNIC	259
	Figure 126: VM with Multiple vNICs Shown in the VM Tree	260
	Figure 127: Policy Per vNIC	261
	Figure 128: Applying Policy to Individual vNICs	263
	Figure 129: Creating a VM Admin Administrator Account	271
	Figure 130: Adding a New Administrator	272
	Figure 131: Changing the Global Administrator Password	273
	Figure 132: Global Administrator Changing the Password of Another Administrator	274
	Figure 133: Administrators Changing Their Password	275
	Figure 134: Configuring Machines Information	279
	Figure 135: Syslog Entry Including VM Name and Log Tag	279
Chapter 16	vGW Series Security Settings	281
	Figure 136: Global Settings Page	283
	Figure 137: Changing the vGW Security VM Management Interface IP Address	287
	Figure 138: Configuring Network Monitoring for Individual vGW Security VMs	290
	Figure 139: IDS Settings Page	293
	Figure 140: IDS Updates Pane	294
	Figure 141: Signatures in a Signature Group	296
	Figure 142: Signature Details	297
	Figure 143: Configuring a Smart Group As a Policy Group	304
	Figure 144: Configuring Policy Rules for a Smart Group with Policy Group Enabled	305
	Figure 145: Creating a Smart Group Using Basic Mode	308
	Figure 146: The Smart Group Editor in Advanced Mode Using Regular Expressions	310
Chapter 17	vGW Security Design Appliance Settings	319
	Figure 147: Network Configuration Settings	324

	Figure 148: Settings Module Backup and Restore Settings	330
	Figure 149: Settings Module Backup and Restore Settings	331
Chapter 19	High Availability and Fault Tolerance	343
	Figure 150: Configuring Network Monitoring and NetFlow Settings	348
	Figure 151: Configuring the Secondary vGW Security Design VM	351
Part 3	Juniper Networks Products Interoperability	
Chapter 20	vGW Series Interoperability with Juniper Networks Products	359
	Figure 152: vGW Series Configuration for Syslog and NetFlow to a STRM Series Device	365
	Figure 153: STRM Source Log Definition for vGW Series	366

List of Tables

Part 1	vGW Series Overview	
Chapter 4	vGW Series IPv6 Support	49
	Table 1: IPv6 Basic Packet Header Fields	51
	Table 2: IPv6 Extension Headers	52
Part 2	vGW Series Configuration and Management	
Chapter 5	The vGW Security Design VM and the vGW Security VM	67
	Table 3: Taskbar Icons	71
	Table 4: Virtual Machine State Icons	75
Chapter 6	vGW Series Main Module	77
	Table 5: vGW Series Status Icons	79
Chapter 7	vGW Series Network Module	85
	Table 6: Using Advanced Options for Filtering Network Data	89
Chapter 8	vGW Series Firewall Module	93
	Table 7: Firewall Policy Configuration Settings	103
	Table 8: Firewall Policy Icons	105
Chapter 11	vGW Series Introspection Module	163
	Table 9: Add Enforcer Profile: Selecting the Gold Image and VMs to Be Compared Against It	174
	Table 10: Edit Enforcer Profile Options	174
	Table 11: VM Gold Image Allowed Deviations	174
	Table 12: Actions	175
	Table 13: Compliance Rule Specifications	175
	Table 14: Scan Definition Options	176
Chapter 12	vGW Series Compliance Module	183
	Table 15: Compliance Rule Creation Parameters	186
Chapter 15	vGW Series Application Settings	205
	Table 16: Smart Group Attributes for vNICs When Policy per vNIC Is Enabled	266
	Table 17: vGW Series Built-In Administrator User Types	270
Chapter 16	vGW Series Security Settings	281
	Table 18: Operators for Creating Smart Groups Using Regular Expression	310
	Table 19: Smart Group Attributes	311

About This Guide

- vGW Series Documentation on page xxiii
- Objectives on page xxiii
- Audience on page xxiii
- Documentation Conventions on page xxiv
- Obtaining Documentation on page xxiv
- Documentation Feedback on page xxiv
- Requesting Technical Support on page xxiv

vGW Series Documentation

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/>.

Objectives

The Juniper Networks vGW Series for virtualized environments runs as integrated software on VMware ESX/ESXi hosts. This guide applies to vGW Series running in the VMware environment. It explains:

- How to integrate the vGW Series with the VMware environment, and install, configure, and administer it.
- How to use the vGW Security Design VM modules that give you advanced methods to secure, administer, and monitor your virtualized environment.

Audience

This guide is intended for the following audiences:

- Virtual infrastructure administrators

This group includes cloud service providers who offer virtual machines and other virtualized resources to enterprises and other business organizations.

These administrators manage virtualized servers, hypervisors, virtual machines, and other virtualized infrastructure. In the cloud deployment, the administrator also manages the physical resources on which virtualized servers and virtual machines run.

- Enterprise network and security administrators

This group includes enterprise personnel who are responsible for network and security infrastructure.

Presently these administrators might manage a virtualized data center, or they might be consolidating and re-architecting their data center for virtualization. They might deploy virtualized services and products within their data center using a private cloud or a hybrid cloud.

Documentation Conventions

Because the vGW Series product is largely presented through a graphical user interface, this document describes the icons, symbols, and navigation methods.

Obtaining Documentation

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

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document name
- Page number
- Software release version

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: <https://www.juniper.net/alerts/>
- 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

vGW Series Overview

This part includes the following chapters:

- [Introduction to vGW Series on page 3](#)
- [vGW Series Resource Requirements on page 15](#)
- [vGW Series Integration with the VMware Environment on page 21](#)
- [vGW Series IPv6 Support on page 49](#)

CHAPTER 1

Introduction to vGW Series

This chapter includes the following topics:

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Architecture on page 4](#)
- [vGW Security Design VM Modules \(VMware\) on page 6](#)
- [Understanding Cloud Computing and vGW Series on page 11](#)
- [Understanding Hypervisors and vGW Series on page 12](#)
- [Understanding the VMware Infrastructure and vGW Series on page 13](#)

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 11](#)
 - [Understanding Hypervisors and vGW Series on page 12](#)
 - [Understanding the vGW Security Design VM on page 67.](#)
 - [Understanding the vGW Security VM on page 68](#)

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.

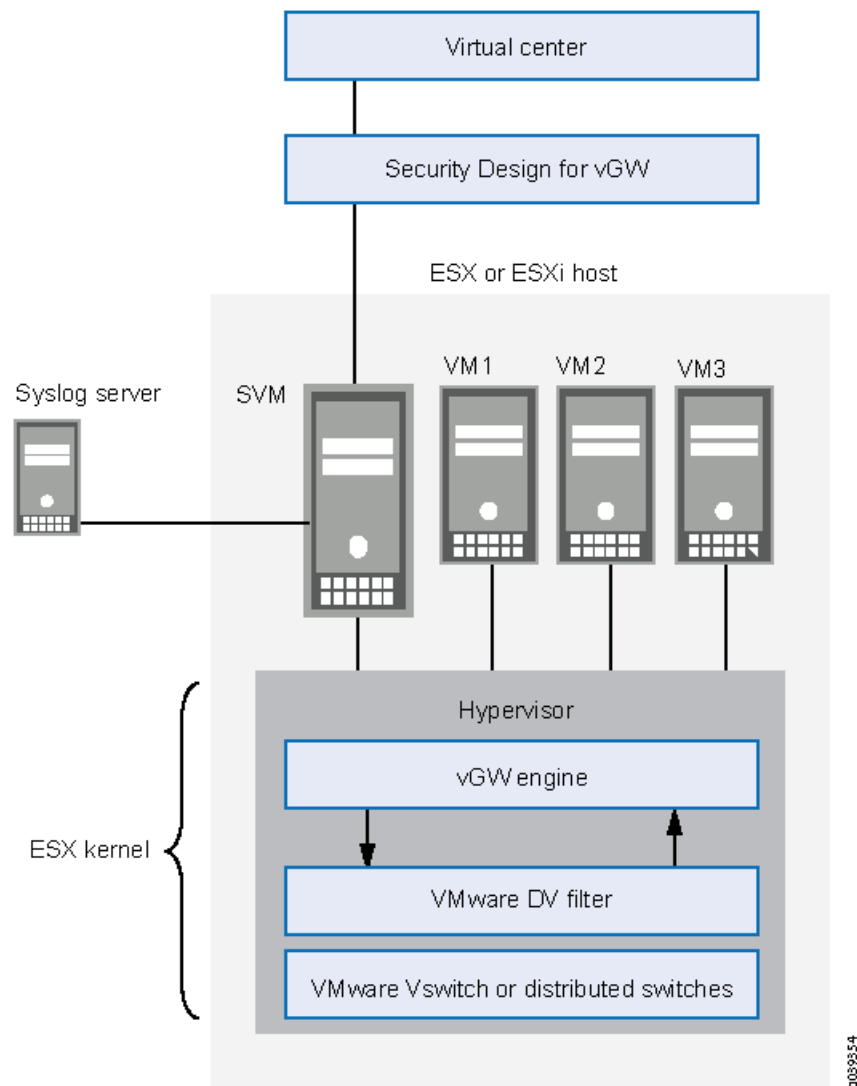
In the VMware virtualized environment, port groups are used to aggregate multiple ports under a common configuration. They serve as an anchor point for virtual machines that connect to labeled networks.

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 on page 364](#)
- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 349](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)
- [Integrating the vGW Series with VMware Using the Settings Module on page 209](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 21](#)
- [Understanding Hypervisors and vGW Series on page 12](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Series High Availability Solution on page 343](#)

- [Understanding the vGW Series Hypervisor and Extended VM Security on page 188](#)
- [Understanding vGW Series Fault Tolerance Support on page 354](#)
- [Viewing the vGW Series Logs on page 336](#)

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.

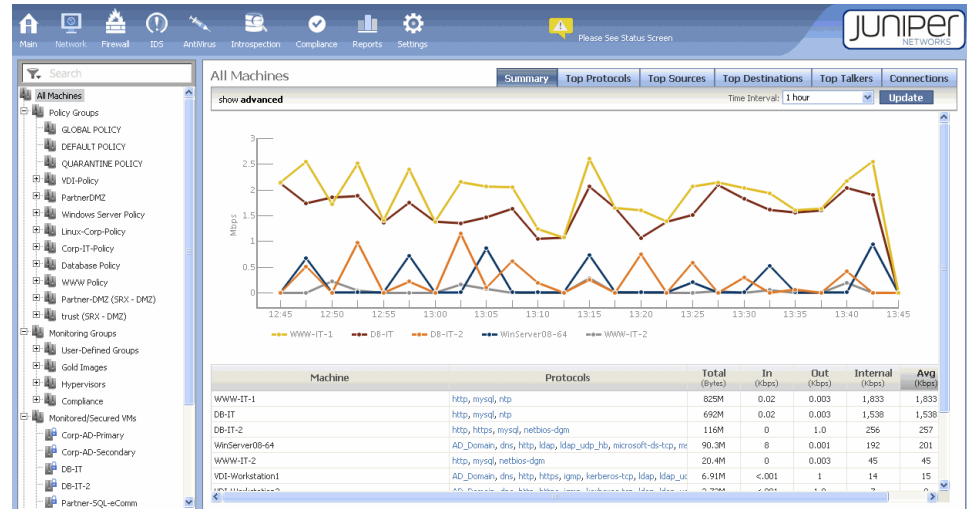
- Main. See “Understanding the vGW Series Main Module” on page 77 and Figure 2 on page 7.

Figure 2: Main Module



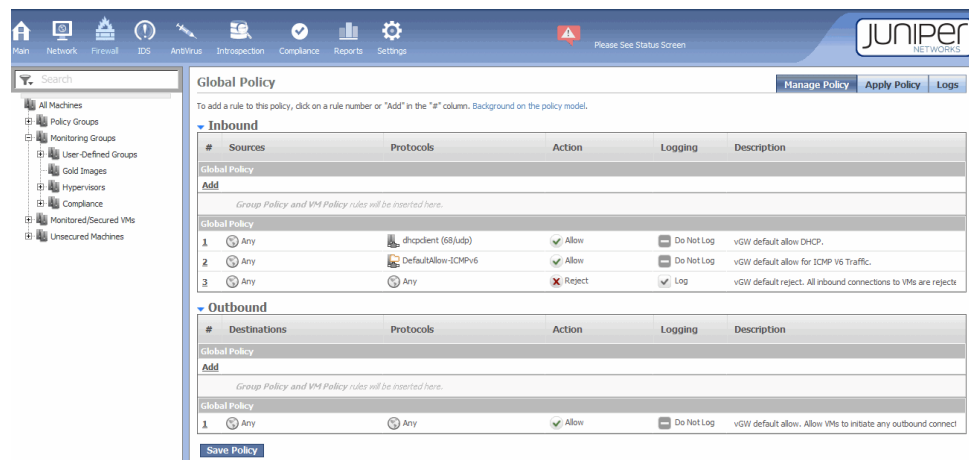
- Network. See “Understanding the vGW Series Network Module” on page 85 and Figure 3 on page 7.

Figure 3: Network Module



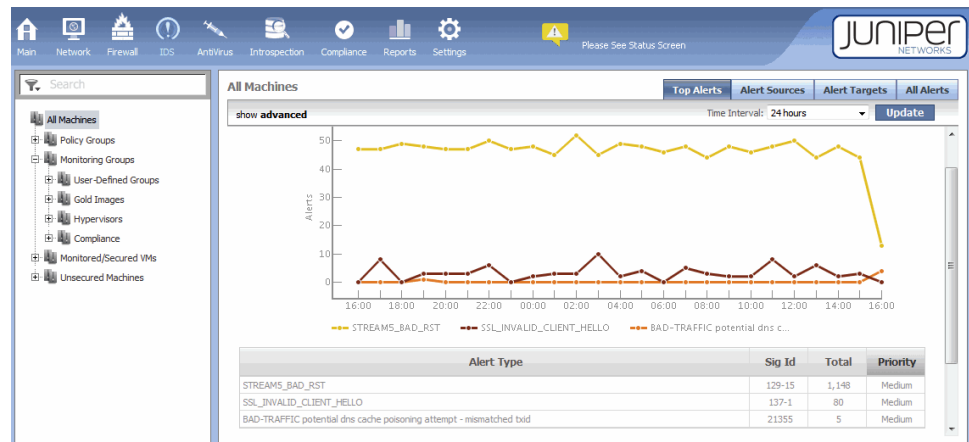
- Firewall. See “Understanding the vGW Series Firewall Module” on page 93 . See Figure 4 on page 8.

Figure 4: Firewall Module



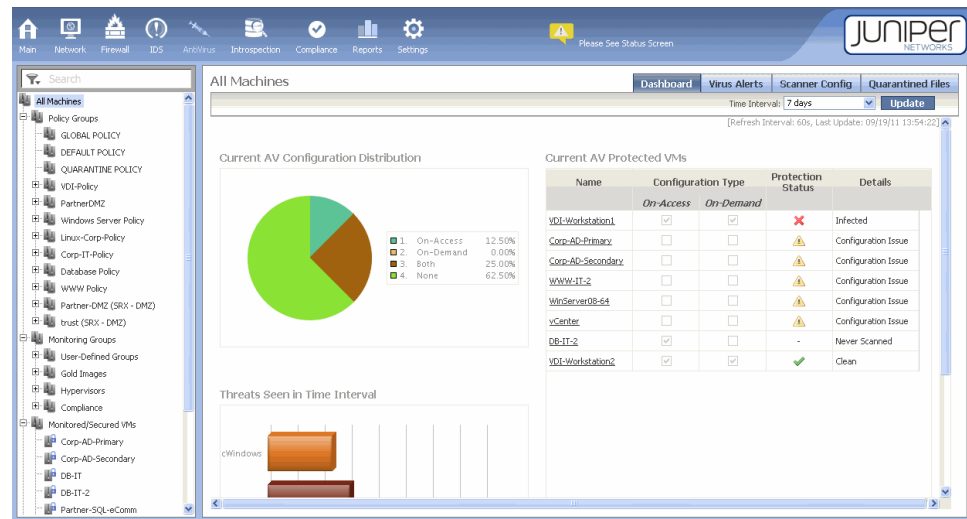
- IDS. See "Understanding the vGW Series IDS Module" on page 121 and Figure 5 on page 8.

Figure 5: IDS Module



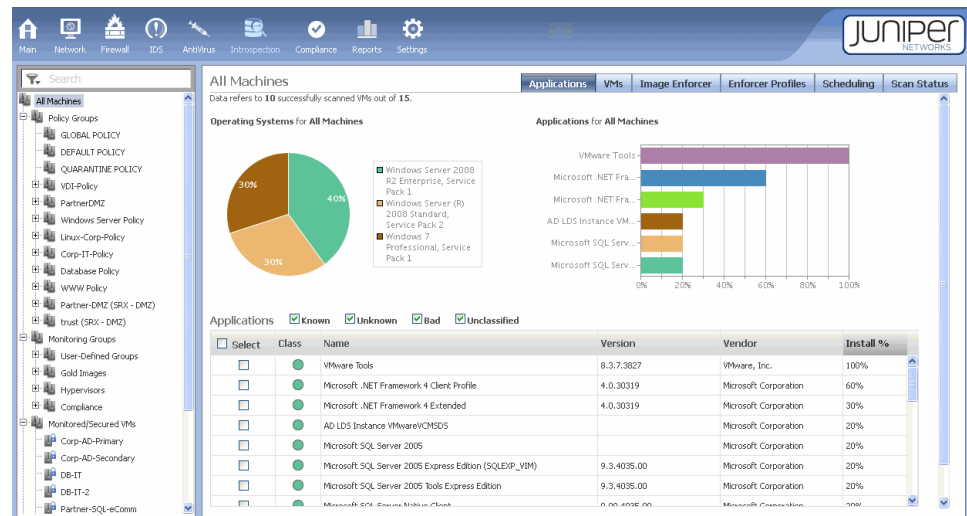
- AntiVirus. See "Understanding vGW Series AntiVirus" on page 131 and Figure 6 on page 9.

Figure 6: AntiVirus Module



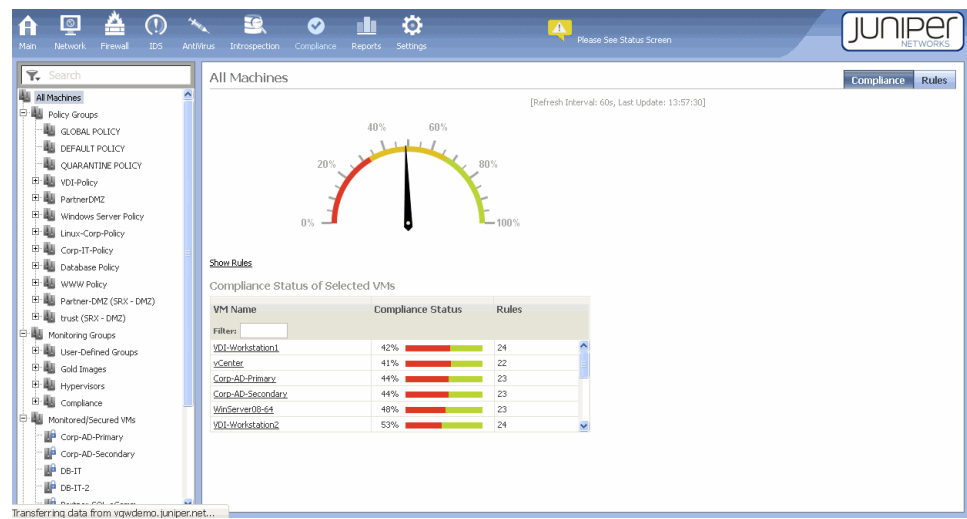
- Introspection. See “Understanding the vGW Series Introspection Module” on page 163 and Figure 7 on page 9.

Figure 7: Introspection Module



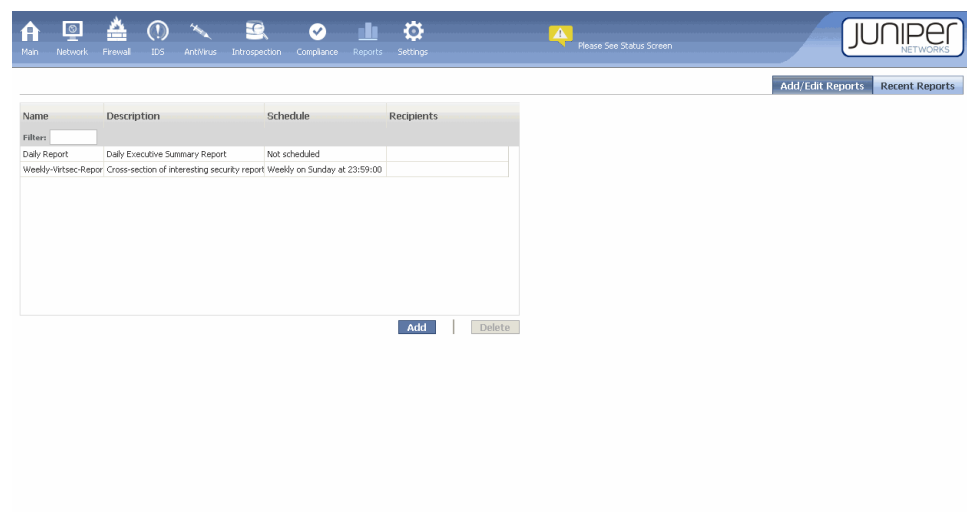
- Compliance. See “Understanding the vGW Series Compliance Module” on page 183 and Figure 8 on page 10.

Figure 8: Compliance Module



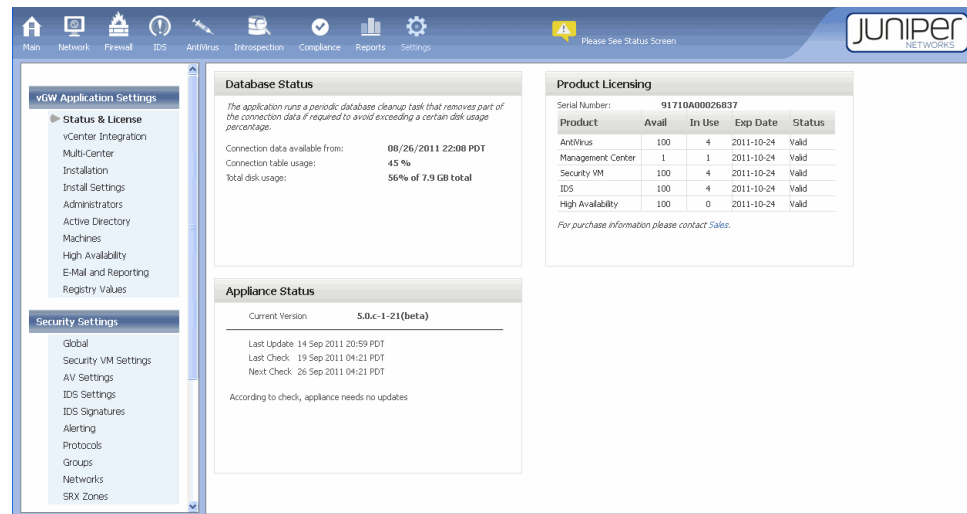
- Reports. See “Understanding the vGW Series Reports Module” on page 193 and Figure 9 on page 10.

Figure 9: Reports Module



- Settings. See “Understanding the vGW Series Settings Module” on page 203 and Figure 10 on page 11.

Figure 10: Settings Module



Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding Cloud Computing and vGW Series on page 11](#)
- [Understanding Hypervisors and vGW Series on page 12](#)
- [Understanding the vGW Security VM on page 68](#)

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 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 11](#)
- [Understanding vGW Series IPv6 Support on page 54](#)

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 13](#)
- [Understanding VMware ESX and ESXi Hosts and the vGW Series on page 13](#)
- [Understanding VMotion and vGW Series on page 14](#)

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 23](#)
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 15](#)
- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 24](#)
- [Integrating the vGW Series with VMware Using the Settings Module on page 209](#)
- [Understanding the vGW Series High Availability Solution on page 343](#)
- [Understanding vGW Series Fault Tolerance Support on page 354](#)
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 15](#)

CHAPTER 2

vGW Series Resource Requirements

This chapter covers the prerequisites and resources required to import the vGW Series into the VMware environment, and install and run it. It includes the following section:

- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 15](#)

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 15](#)
- [Virtual Appliance System Requirements on page 16](#)
- [vGW Series VMware vSwitch Requirements on page 17](#)
- [VMware Port Group Requirements on page 18](#)
- [Virtualized NIC Requirements on page 18](#)

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 vmsservice-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_  
vmsservice-altor  
vmsservice-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” on page 261](#).

**Related
Documentation**

- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 24](#)
- [Understanding vGW Series on page 3](#)
- [Integrating the vGW Series with VMware Using the Settings Module on page 209](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)

CHAPTER 3

vGW Series Integration with the VMware Environment

This chapter explains how to integrate the vGW Series with the VMware infrastructure. It also explains the two VMware VMsafe installation modes that are used to install the product. In both modes, the vGW Security Design VM automatically makes configuration changes to the VMware vCenter. This automatic process lowers administrative complexity and reduces configuration errors. You can use the vGW Security Design VM to automatically remove these configuration changes later if necessary.

This chapter includes the following sections:

- [Preparing to Integrate the vGW Series with the VMware Environment on page 21](#)
- [vGW Series VMsafe Firewall + Monitoring and VMsafe Monitoring Modes on page 22](#)
- [Understanding vGW Series Environment Time Synchronization on page 23](#)
- [Understanding the Open Virtualization Format OVA Template Method on page 23](#)
- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 24](#)
- [Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware on page 34](#)
- [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 36](#)
- [Setting Up vGW Series on page 36](#)

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 13](#)
- [Understanding vGW Series on page 3](#)
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 15](#)

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](#)" on [page 300](#).

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 68](#)
 - [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
 - [Understanding the vGW Security Design VM on page 67](#)

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” on page 328.

- Related Documentation**
- [Understanding vGW Series on page 3](#)
 - [Understanding the vGW Security Design VM on page 67](#)
 - [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 349](#)
 - [Installing a Secondary vGW Security VM for High Availability on page 353](#)

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”](#) on page 349.

Related Documentation

- [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 36](#)
- [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 24](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 21](#)
- [Understanding the VMware Infrastructure and vGW Series on page 13](#)
- [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 15.

This topic includes the following sections:

- [Requirements on page 24](#)
- [Overview on page 25](#)
- [Downloading the vGW Series OVA Combo Package on page 25](#)
- [Integrating the vGW Series with the VMware Infrastructure on page 26](#)

Requirements

For information, see [“vGW Series Prerequisites and Resource Requirements for the VMware Environment”](#) on page 15.

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 34.](#)
- For details, see [“Using the OVA Single File Method to Integrate the vGW Security VM with VMware” on page 36.](#)

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:

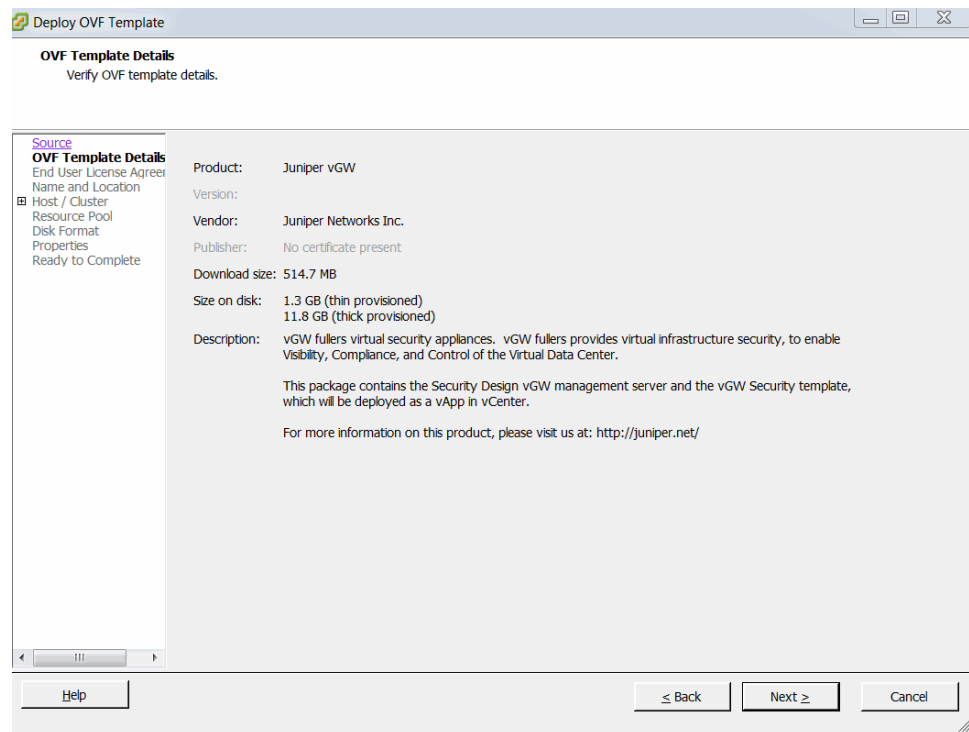
1. Using the vSphere client, load the bundled OVA file. Select **Deploy OVF Template** from the File menu.
2. Enter the download filename or its URL in the Deploy from file or URL box—for example, enter: `c:\temp\vGW_Combo_5.5_#-#-#-#-#.ova`—and click **Next**.

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**.

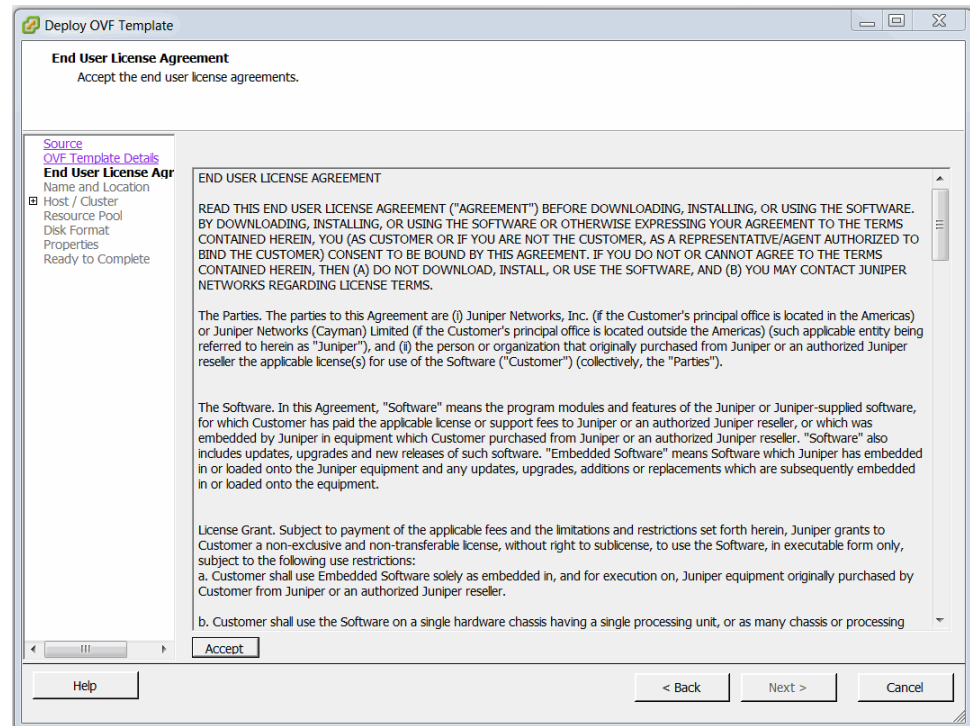
Before the wizard unbundles the OVA package, verify that it contains the vGW Series appliances. The OVA template summary also specifies the disk space requirements for thick and thin provisioning. See [Figure 11 on page 26](#).

Figure 11: OVA Template Details Page



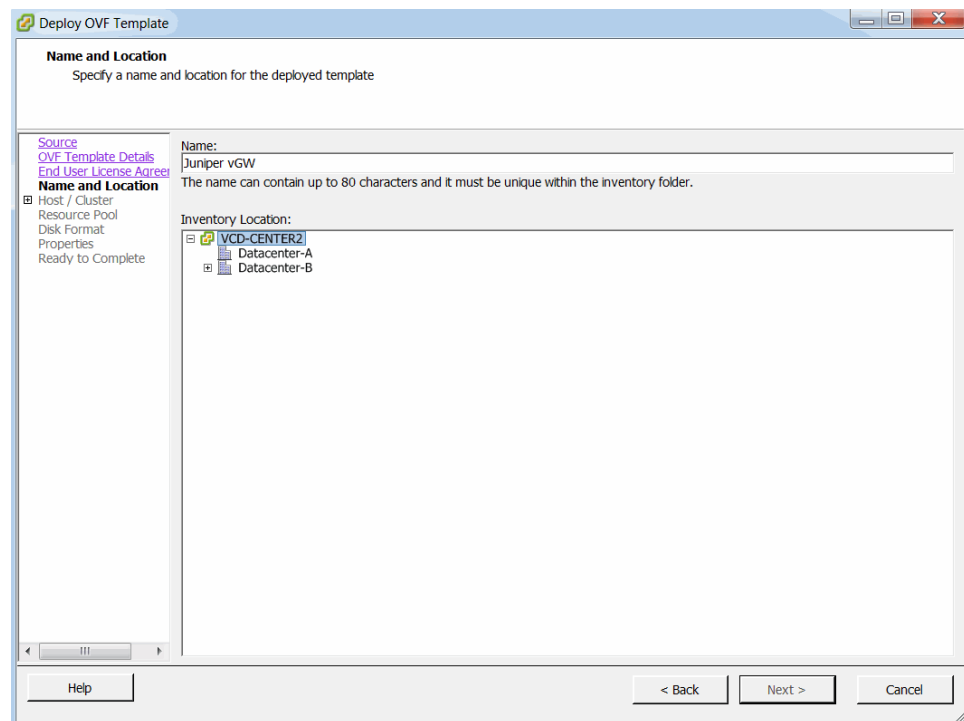
4. Accept the vGW Series license agreement, and click **Next**. See [Figure 12 on page 27](#).

Figure 12: OVA File Deployment License Agreement




5. Specify a name for the vApp that will be created and a storage location. See [Figure 13 on page 28](#).

Figure 13: Naming the vApp



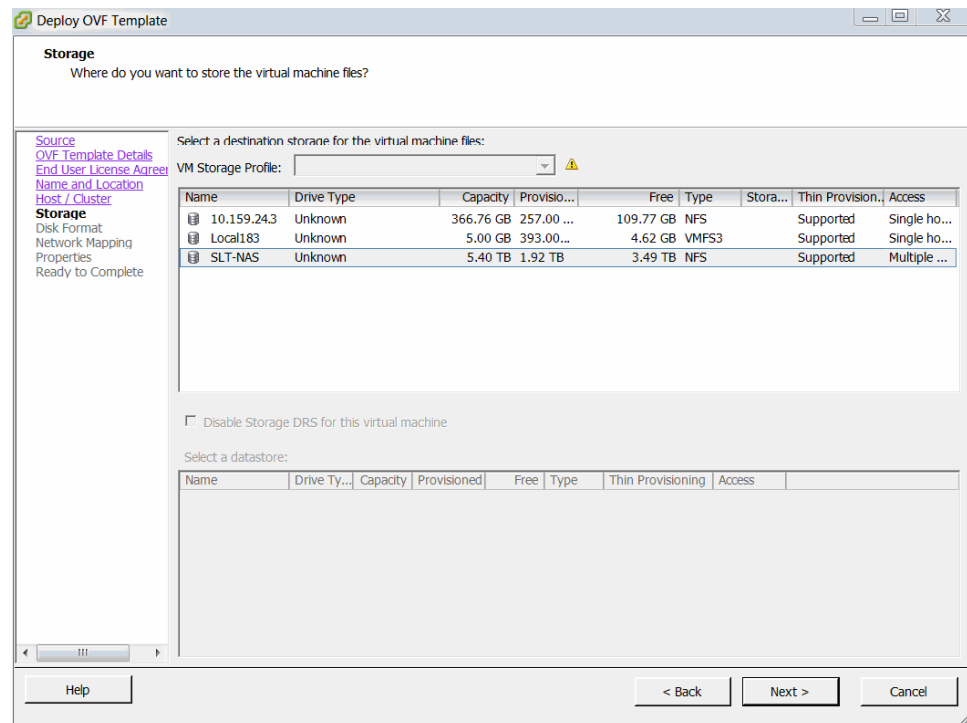
6. Specify the host or host/cluster on which to run the deployed template. We recommend that you use a network storage device (NAS) so that it can be migrated through VMotion for space optimization. See [Figure 14 on page 29](#).

Figure 14: Specifying the Host and Cluster



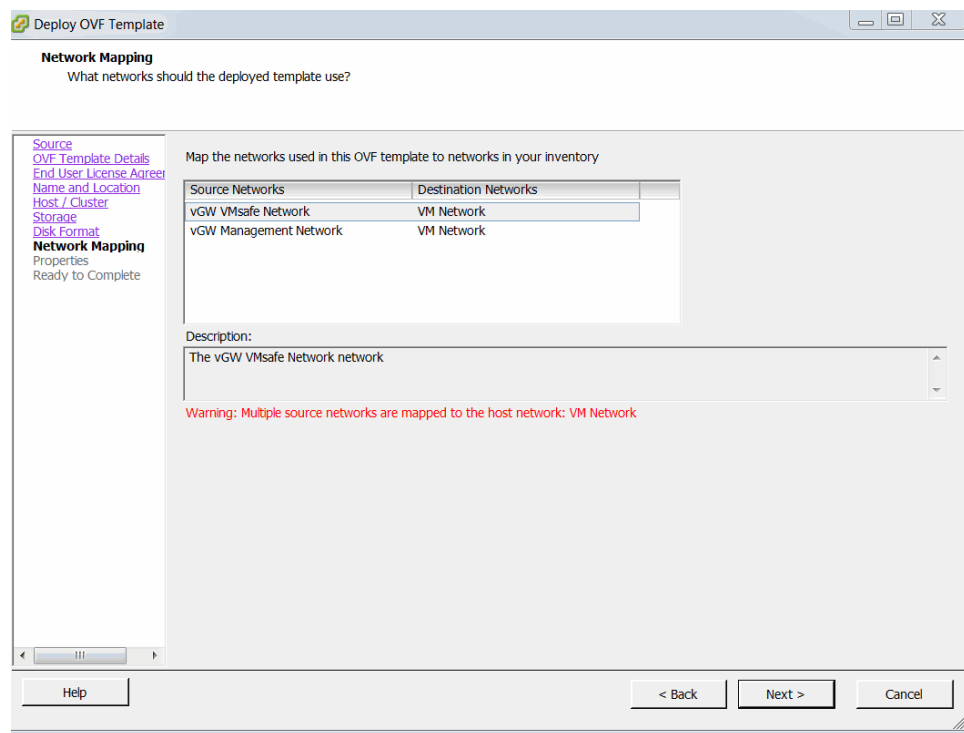
7. Select the datastore. Do not use a read-only datastore. [Figure 15 on page 29.](#)

Figure 15: Selecting the Storage



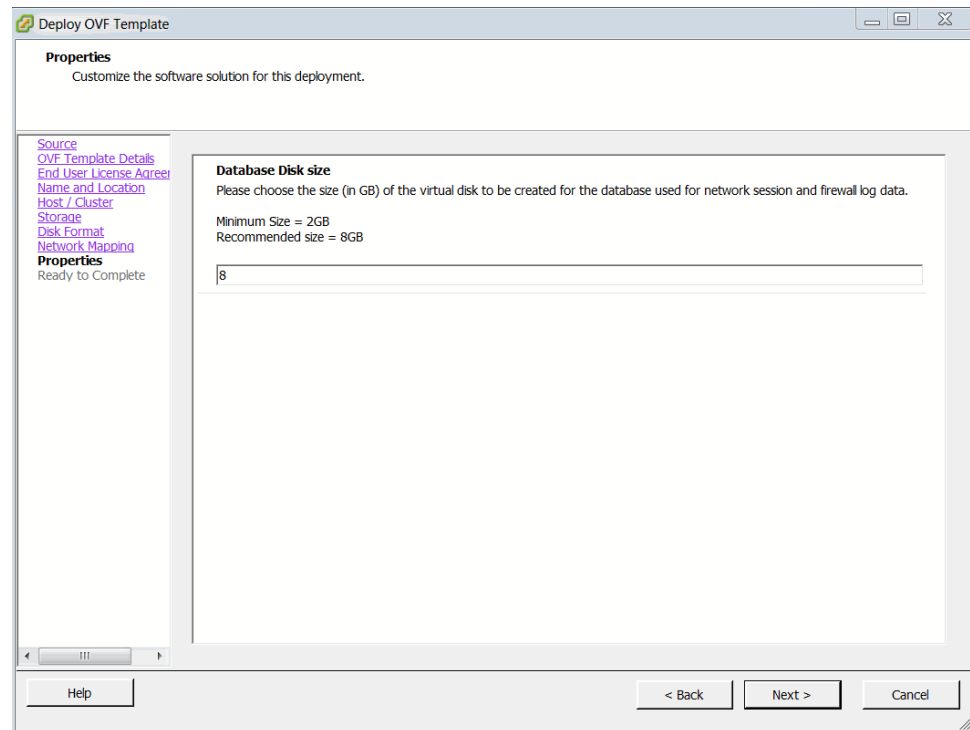
8. Select the disk format. Accept the thick provisioned format default. Thick provisioning preallocates all required space for the product.
9. Map the networks. Set the vGW management network to a destination network that is accessible to vCenter and the vGW Security Design VM. See [Figure 16 on page 30](#).

Figure 16: 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 17 on page 31](#).

Figure 17: Specifying the Database Disk Size




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 18 on page 31](#).

Figure 18: Verifying That the Configuration Is Correct




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 19 on page 32](#).

Figure 19: 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:

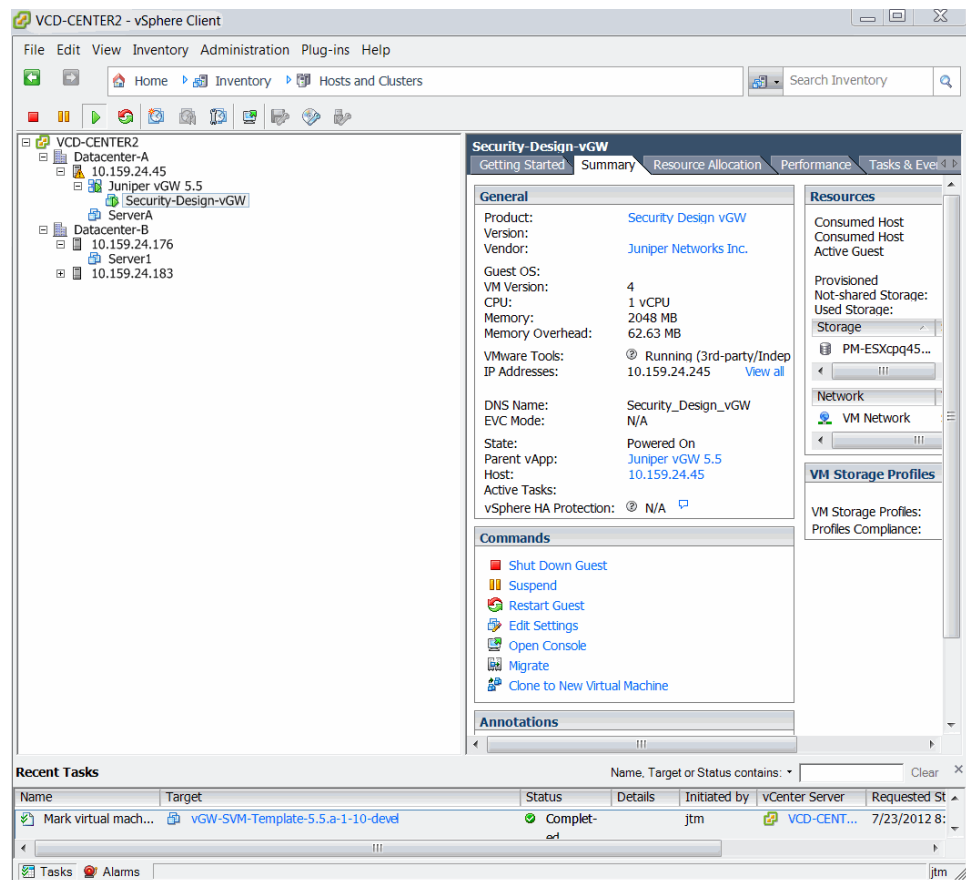
- a. Convert the vGW-SVM-Template VM to a template that the vGW Security Design VM and installer can use to instantiate a vGW Security VM on each ESX/ESXi host to be secured.

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

- b. Right-click the vGW Security Design VM and power it on. [Figure 20 on page 34](#) shows the vCenter summary information for the vGW Security Design VM.

[Figure 20 on page 34](#) shows the vCenter summary information for the vGW Security Design VM.

Figure 20: vGW Security Design VM Summary Tab in vCenter



Related Documentation

- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 15](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the VMware Infrastructure and vGW Series on page 13](#)
- [Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware on page 34](#)
- [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 36](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security VM on page 68](#)

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 15](#)
 - [Understanding vGW Series on page 3](#)
 - [Understanding the VMware Infrastructure and vGW Series on page 13](#)
 - [Using the OVA Single File Method to Integrate the vGW Security VM with VMware on page 36](#)
 - [Understanding the vGW Security Design VM on page 67](#)
 - [Understanding the vGW Security VM on page 68](#)

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 15](#)
 - [Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware on page 34.](#)
 - [Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure on page 24](#)
 - [Preparing to Integrate the vGW Series with the VMware Environment on page 21](#)

Setting Up vGW Series

After you download and integrate vGW Series with the VMware environment and power on the vGW Security Design VM, you can configure its basic system parameters. This

topic explains how to connect to the vGW Series Security Design VM to configure basic settings. It describes how to use the vGW Series wizard to configure those settings initially.

This topic includes the following sections:

- [Determining the vGW Security Design VM's Default IP Address on page 37](#)
- [Changing or Setting the IP Address for the vGW Security Design VM on page 38](#)
- [Connecting to the vGW Security Design VM and Configuring Basic Settings on page 39](#)

Determining the vGW Security Design VM's Default IP Address

To access the vGW Security Design VM, you enter its IP address in a supported Web browser.

When you powered on the vGW Security Design VM during vGW Series integration with VMware, which is described in [“Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure” on page 24](#), it acquired an IP address that you can view on the vCenter Summary page.

By default, the vGW Security Design VM is configured to use IPv4 DHCP to acquire its address. If problems occur and it cannot obtain an IP address in this manner, it tries other methods. In order, these are the three methods that the vGW Security Design VM uses in an attempt to obtain an IP address:

- IPv4 DHCP
- IPv6 autoconfiguration

With stateless IPv6 autoconfiguration, the vGW Security Design VM acquires its IPv6 address automatically without the intervention of a DHCP server.

- IPv6 DHCPv6



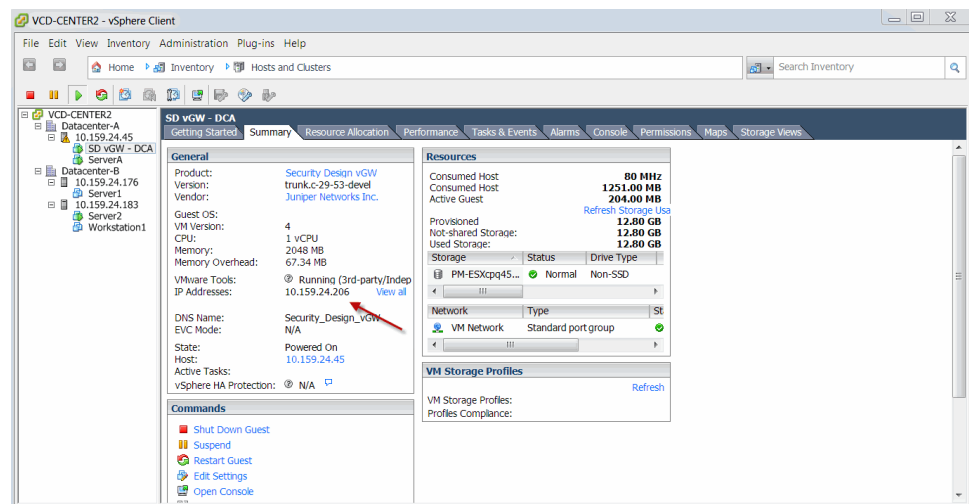
NOTE: If DHCP is not available on the vGW Security Design VM network, you can log into the console using **admin** for both the username and the password. Type **config network** at the command prompt and proceed through the options to assign an IP address. After an IP address is set—whether DHCP or static, you can access the vGW Security Design VM through a Web browser.

To view the IP address bound to the vGW Security Design VM:

1. Launch the VMware vSphere Client, and select the vGW Security Design VM icon on the left navigation pane.
2. Select the **Summary** tab.

The IP address that was acquired appears in the IP Addresses: field. See [Figure 21 on page 38](#).

Figure 21: Viewing the vGW Security Design VM IP Address in VMware



After you have obtained the IP address, follow the instructions in “[Connecting to the vGW Security Design VM and Configuring Basic Settings](#)” on page 39 to configure the basic settings.

Changing or Setting the IP Address for the vGW Security Design VM


You can use the vGW Series command-line interface (CLI) to set the IP address for the vGW Security Design VM.

To use the vGW Series CLI from the vCenter console:

1. Launch the VMware vSphere Client.
2. Right-click the vGW Security Design VM icon on the left navigation panel to display a list of options.
3. Select the third option on the list, **Open Console**. Alternatively you can select the **Console** tab, as shown in [Figure 22 on page 38](#).

The console window appears.

Figure 22: vGW Security Design VM IP Addresses on the vGW Series CLI Console



4. At the CLI prompt enter **config network**. Enter **admin** for both the username and password. In this mode you can configure the IP address for the vGW Security Design

VM. You can specify an IP address for either IP protocol family. See [Figure 23 on page 39](#).

Figure 23: Configuring an IP Address for the vGW Security Design VM

```
(Cmd) config network
1) ipv4 dhcp          4) ipv6 autoconfiguration
2) ipv4 static        5) ipv6 static
3) ipv6 dhcp
What type of IP address do you want to set ?
```

- In response to the prompt **What type of IP address do you want to set?** enter the number preceding the type of IP address that you want to be assigned to the vGW Security Design VM and how it is to be acquired.

For example, the administrator might enter **4** for **4) ipv6 autoconfiguration**.

- The vGW CLI gives you the opportunity to cancel by presenting the prompt **Are you sure?**. If you enter **y** for **y(es)**, vGW Series shuts down the interface and brings it back up with the new IP address.

Connecting to the vGW Security Design VM and Configuring Basic Settings

This section explains how to set up the vGW Security Design VM initially.

- Using a supported Web browser, connect to the vGW Security Design VM management interface through HTTPS. Enter the IP address of the vGW Security Design VM in the Web browser.

This is the IP address that was assigned when you powered on the vGW Security Design VM.

vGW Series supports the following Web browsers:

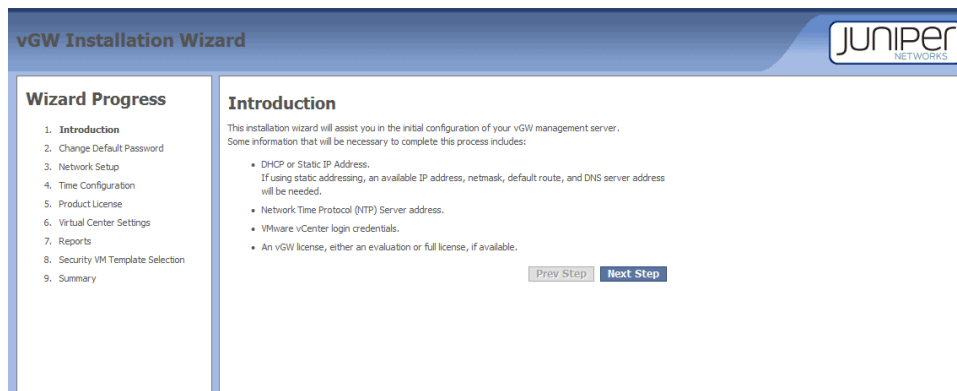
- Microsoft Internet Explorer 7, 8, and 9
 - Mozilla Firefox 3 or later
- Enter **admin** for both the username and password. See [Figure 24 on page 40](#).

Figure 24: Logging In to the vGW Security Design VM



3. Read the information message, and review the process overview shown in the Wizard Progress pane. See [Figure 25 on page 40](#).

Figure 25: vGW Series Installation Wizard Overview



4. Change the default vGW global administrator account password—admin—that you used to log in.

You must change the default password. See [Figure 26 on page 41](#). Store the new password in a secure location. It is difficult to recover a lost or forgotten global administrator account password. If you wish, you can change the password that you specify here later, but to do so you must enter your current password, which would be the one that you configured here.



TIP: You can integrate administration accounts with the vGW Security Design VM after the installation is completed. For information on how to do that later, see [“Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module” on page 269](#).

Figure 26: Changing the Default Password

vGW Installation Wizard

Wizard Progress

1. Introduction
2. **Change Default Password**
3. Network Setup
4. Time Configuration
5. Product License
6. Virtual Center Settings
7. Reports
8. Security VM Template Selection
9. Summary

Change Default Password

Please select a new password for the default global administrator account. The user-name for this account is "admin". The "admin" account can also be used to login to the console of the virtual machine for advanced configuration.

Current Password:

New Password:

Confirm New Password:

[Prev Step](#) [Next Step](#)

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



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

Set the correct destination network for the vGW Security Design VM and leave the VMsafe Network unchanged. At this point you can configure other network information for the vGW Security Design VM, such as whether to use dual stack for it and how it obtains its management interface addresses.

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.

For its management interface addressing mode, either accept the default dual stack values of DHCP for IPv4 and DHCPv6 for IPv6 or change the values by selecting:

- IPv4
 - DHCP (Default): To obtain an IPv4 address, by default the vGW Security Design VM is configured to use DHCP. You do not need to specify additional information.
 - Static IP. If you select **Static IP**, you must specify a static IPv4 address and its network mask routing prefix, and the default gateway to assign to the vGW Security Design VM.
- IPv6
 - DHCPv6 (Default): To obtain an IPv6 address, by default the vGW Security Design VM is configured to use DHCPv6. You do not need to specify additional information.
 - Autoconfiguration. If you select **Autoconfiguration**, stateless address autoconfiguration is used to obtain the IPv6 address. It 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.
 - Static IP. If you select **Static IP**, you must specify a static IPv6 address, including the IPv6 address prefix (the initial bits of the address that denote the network address, akin to a netmask), and the default gateway to use for it.



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.

If you do not want the vGW Security Design VM to be configured for dual stack which is its default configuration, you can change the configuration in the following way:

- To use only IPv4 for vGW Security Design VM management communication with its vGW Security VMs, disable IPv6. On the displayed list for the IPv6: box, select **Disabled**.
- To use only IPv6 for vGW Security Design VM management communication with its vGW Security VMs, disable IPv4. On the displayed list for the IPv4: box, select **Disabled**.

How you configure addressing for the vGW Security Design VM management center affects its communication with its vGW Security VMs in the following way:

- In an environment in which both the vGW Security Design VM and the vGW Security VM are configured for dual stack, communication problems between the vGW Security Design VM management interface and that of the vGW Security VMs should not occur.
- In an environment in which the vGW Security Design VM is configured for dual stack but one or more of the vGW Security VMs is not, communication problems between their management interfaces should not occur.
- In an environment in which the vGW Security Design VM is not configured for dual stack but all of the vGW Security VMs are, communication problems between their management interfaces should not occur.
- In an environment in which neither the vGW Security Design VM nor one or more vGW Security VM is configured for dual stack, in any case in which the IP address type of the management interfaces of the vGW Security Design VM and the vGW Security VM differ—one might belong to the IPv6 protocol family and the other to the IPv4 protocol family—communication problems will occur. The vGW Security Design VM will not be able to connect to the vGW Security VM to carry out any procedures.

You can make these changes during the installation process, as shown [Figure 27 on page 43](#), or you can make them later, after you complete the initial configuration.

Figure 27: Configuring Network Settings for the vGW Security Design VM

In the latter case, you use the Settings module Appliance Settings > Network Settings page, which is the same page shown in [Figure 27 on page 43](#), only it is arrived at differently. For additional details, see “[Configuring the vGW Series Network Settings](#)” on page 323.



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

6. Set the system time.

Set the correct time zone, and then specify the NTP servers for your environment. See [Figure 28 on page 44](#).

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 you do not have an internal NTP server, then you must clear all entries shown in this window and set the time manually.

To do this, you use the vGW Series CLI that you run from the vCenter console.

Figure 28: Configuring the Time Server

The screenshot shows the vGW Installation Wizard interface. On the left, the 'Wizard Progress' sidebar lists steps 1 through 9, with step 4, 'Time Configuration', highlighted. The main panel has a title bar 'vGW Installation Wizard' and the Juniper Networks logo. Below the title bar, a message states: 'To ensure accurate reporting of network activity, it is important to configure a time server to be used as the source for the vGW management server.' The 'Time Configuration' section includes a 'Select Timezone' dropdown menu set to 'US/Pacific'. Below this are four input fields for NTP servers, each with a default value: '0.altor.pool.ntp.org', '1.altor.pool.ntp.org', '2.altor.pool.ntp.org', and an empty field for the fourth. At the bottom of this section are 'Prev Step' and 'Next Step' buttons.

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

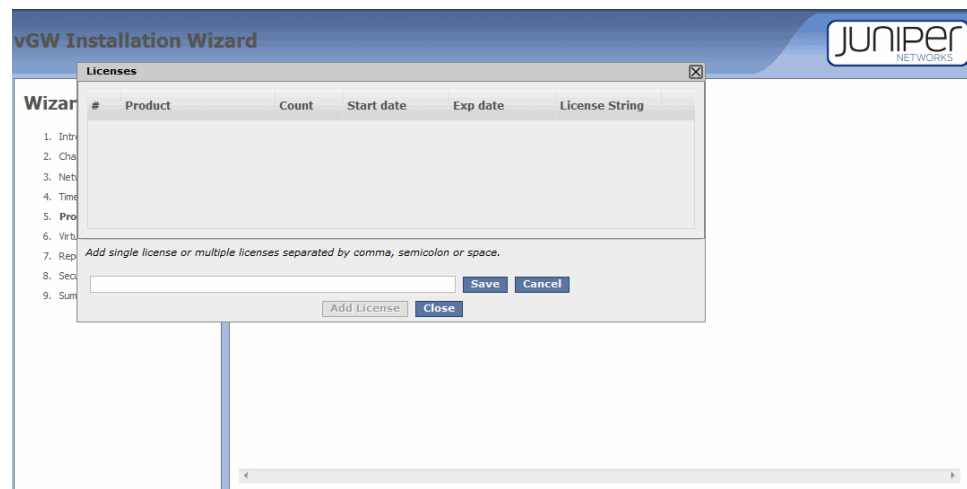
Next, the wizard prompts you for license information, as shown in [Figure 29 on page 44](#).

Figure 29: Configuring vGW Series Product Licensing

The screenshot shows the vGW Installation Wizard interface. On the left, the 'Wizard Progress' sidebar lists steps 1 through 9, with step 5, 'Product License', highlighted. The main panel has a title bar 'vGW Installation Wizard' and the Juniper Networks logo. Below the title bar, a message states: 'Enter permanent license or continue in evaluation mode.' The 'Product Licensing' section includes a paragraph explaining the 30-day evaluation mode and the option to enter a permanent license. Below this is a text input field labeled 'Enter Permanent or long term evaluation License(s)'. At the bottom of this section are 'Prev Step' and 'Continue in 30 Day Eval. Mode' buttons.

If you are using a 30-day evaluation license, you can continue to use vGW Series in that mode. Instead, you can enter one or more permanent licenses or specify a long-term evaluation license. To enter a permanent or extended evaluation license, click **Enter Permanent or long term evaluation License(s)** and enter the license information. See [Figure 30 on page 45](#).

Figure 30: Configuring Permanent Licenses



7. Select the management domain, or scope, for this vGW Security Design VM to manage, and verify that the vGW Security Design VM can establish a connection to vCenter. Then click **Next Step**.

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

- If the connection works properly, a message appears stating that the login was successful, and it identifies 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 exists.

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

To configure a management domain:

- If this vGW Security Design VM will manage all of the vCenter's resources, select **Entire vCenter**.
- If this vGW Security Design VM will participate in a Split-Center configuration, select the data centers or the host clusters for this vGW Security Design VM to manage. To select host clusters, first select the data center that the host clusters belong to.

For information on Split-Center and its configuration options, see [“Understanding the vGW Series Split-Center Feature”](#) on page 240.

Figure 31 on page 46 shows that this vGW Security Design VM is configured to manage two host clusters in Datacenter-B.

Figure 31: Configuring the vGW Security Design VM vCenter Settings

8. (Optional) Configure the e-mail server to use to send reports.

Using this option, you can configure 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 and the content of standard report e-mail. After you configure these parameters, you can test the e-mail connection.

You can also use the Settings module vGW Application Settings > E-Mail and Reporting page to configure this information after the installation is completed.

9. Define a template to use to instantiate vGW Security VMs on ESX/ESXi hosts to secure them.

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 kernel module cannot be loaded or is not present.

You can define whether monitoring is used. 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** selected. You can change this option later if you want to change the behavior for one or more VMs.

10. 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**

- [Understanding the VMware Infrastructure and vGW Series on page 13](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 21](#)
- [Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module on page 269](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 4

vGW Series IPv6 Support

- [Understanding IPv6 Addressing on page 49](#)
- [Understanding vGW Series IPv6 Support on page 54](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments on page 60](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 62](#)
- [Understanding Requirements for Communication Between vGW Series IPv6 Components and Juniper Networks IPv4 Servers on page 63](#)
- [Understanding vGW Series IPv6 Accessibility Requirements for Customer Servers on page 64](#)

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 49](#)
- [IPv6 and IPv4 on page 50](#)
- [IPv6 Address Space, Addressing, and Address Types on page 50](#)
- [The IPv6 Basic Packet Header on page 51](#)
- [The IPv6 Packet Header Extensions on page 52](#)
- [The IPv6 Address Format on page 53](#)
- [Address Assignment and IPv6 on page 54](#)

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 1 on page 51](#).

Table 1: 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.
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.

Table 1: IPv6 Basic Packet Header Fields (*continued*)

Header Name	Bit Length	Purpose
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 2 on page 52](#):



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

Table 2: 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.	

Table 2: IPv6 Extension Headers (*continued*)

Header Name	Purpose
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	<p>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.</p> <p>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.</p>
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.
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	<p>Identifies the host device, or interface on a host, to which the IPv6 packet is to be sent.</p> <p>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.</p>

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 54](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [Understanding vGW Series on page 3](#)

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.

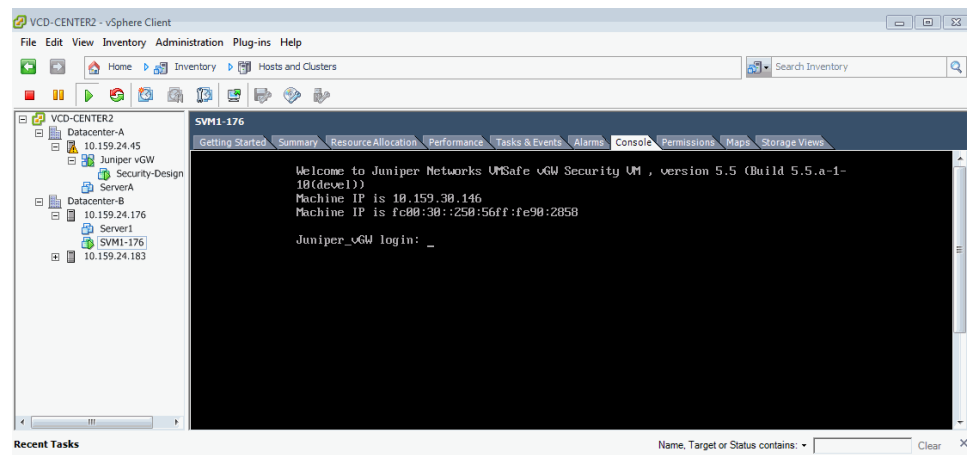
This topic covers how vGW Series displays or allows you to enter IPv6 address in the vGW Security Design VM modules.

- [vGW Security Design VM and vGW CLI Support for IPv6 Addresses on page 55](#)
- [Entering IPv6 Addresses on page 56](#)
- [vGW Series IPv6 Address Representation on page 56](#)
- [IPv4-Mapped IPv6 Addresses on page 57](#)
- [vGW Security Design VM Filter Boxes on page 57](#)
- [Searching the VM Tree for VMs and Hypervisors with IPv6 Addresses on page 57](#)
- [vGW Security Design VM and IPv6 and IPv4 Addressing on page 57](#)
- [vGW Security VM IP Addressing Support on page 57](#)

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 32 on page 55](#) shows the console displaying both the IPv4 and the IPv6 addresses for the SVM-176.

Figure 32: 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” on page 323](#).



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 58](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 62](#)
- [IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments on page 60](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Understanding vGW Series on page 3](#)
- [Configuring vGW Series Installation Settings on page 224](#)

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 58](#)
- [Network Module on page 58](#)
- [Firewall Module on page 58](#)
- [IDS Module on page 59](#)
- [AntiVirus Module on page 59](#)
- [Introspection Module on page 59](#)
- [Compliance Module on page 59](#)
- [Reports Module on page 60](#)
- [Settings Module on page 60](#)

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” on page 106](#).

- 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” on page 299](#).

For details, see the topics that pertain to the Settings module. See [“Understanding the vGW Series Settings Module” on page 203](#).

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series IPv6 Support on page 54](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 62](#)

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 60](#)
- [IPv6 Traffic Handling in Heterogeneous Environments \(with a Mix of vGW Series Component Versions\) on page 61](#)

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 61](#).

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 54](#)
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 62](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [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 62](#)

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 49](#).

Related Documentation

- [Understanding vGW Series IPv6 Support on page 54](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [Understanding vGW Series on page 3](#)

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 \(VMware\) on page 282](#)
- [Understanding vGW Series IPv6 Support on page 54](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [Understanding vGW Series Log Collection on page 333](#)

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 54](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [Understanding Requirements for Communication Between vGW Series IPv6 Components and Juniper Networks IPv4 Servers on page 63](#)
- [Understanding vGW Series on page 3](#)
- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)

PART 2

vGW Series Configuration and Management

This part contains the following chapters:

- The vGW Security Design VM and the vGW Security VM on page 67
- vGW Series Main Module on page 77
- vGW Series Network Module on page 85
- vGW Series Firewall Module on page 93
- vGW Series IDS Module on page 121
- vGW Series AntiVirus Module on page 131
- vGW Series Introspection Module on page 163
- vGW Series Compliance Module on page 183
- vGW Series Reports Module on page 193
- vGW Series Settings Module on page 203
- vGW Series Application Settings on page 205
- vGW Series Security Settings on page 281
- vGW Security Design Appliance Settings on page 319
- vGW Series Status Alerts on page 339
- High Availability and Fault Tolerance on page 343

CHAPTER 5

The vGW Security Design VM and the vGW Security VM

This chapter contains the following topics:

- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding vGW Security Design VM Navigation on page 69](#)
- [Understanding the vGW Series Kernel Module on page 71](#)
- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)

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 33 on page 68](#). The page displays information gathered from the activity of various vGW Security Design VM modules.

Figure 33: Main Module Displayed at Login



Related Documentation

- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)
- [Overview of IPv6 Implementation in the vGW Security Design VM Modules on page 58](#)
- [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 on page 286](#)
- [Understanding the vGW Series Firewall Module on page 93](#)

- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)

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 34 on page 69](#) and [Figure 35 on page 70](#).

You can use the search field to filter VMs in various ways. See [“About the vGW Security Design VM Tree” on page 73](#).

Figure 34: vGW Security Design VM Taskbar


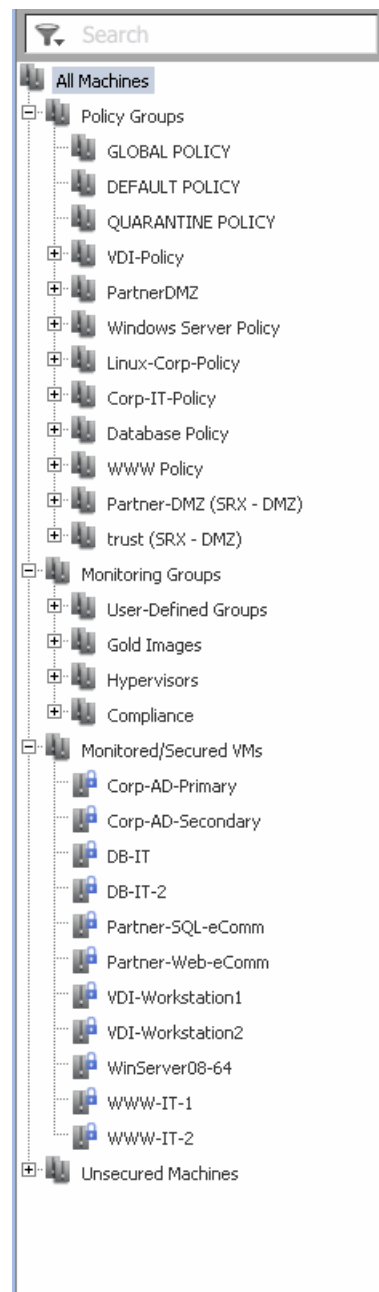


Figure 35: VM Tree



Related Documentation

- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)

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 68](#)
 - [Understanding vGW Series on page 3](#)
 - [Understanding the vGW Security Design VM on page 67](#)

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 36 on page 71](#) shows the taskbar.

Figure 36: vGW Security Design VM Taskbar



[Table 3 on page 71](#) 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 3: 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 77

Table 3: Taskbar Icons (*continued*)

Icon	Module	Description	For details, see:
 Network	Network	Displays a network activity summary, top protocols, sources, destinations, talkers, and connections.	"Understanding the vGW Series Network Module" on page 85
 Firewall	Firewall	Manages and installs policies, and displays logs.	"Understanding the vGW Series Firewall Module" on page 93
 IDS	IDS	Monitors all network traffic or a selected subset of VMs or protocols.	"Understanding the vGW Series IDS Module" on page 121
 AntiVirus	AntiVirus	Protects VMs by detecting malware, and it quarantines affected files or VMs.	"Understanding vGW Series AntiVirus" on page 131
 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.	"Understanding the vGW Series Introspection Module" on page 163
 Compliance	Compliance	Monitors the virtual infrastructure against a predefined set of rules to guarantee all components are configured securely.	"Understanding the vGW Series Compliance Module" on page 183.
 Reports	Reports	Produces detailed system and security reports.	"Understanding the vGW Series Reports Module" on page 193.
 Settings	Settings	Controls configuration settings, including passwords.	"Understanding the vGW Series Settings Module" on page 203

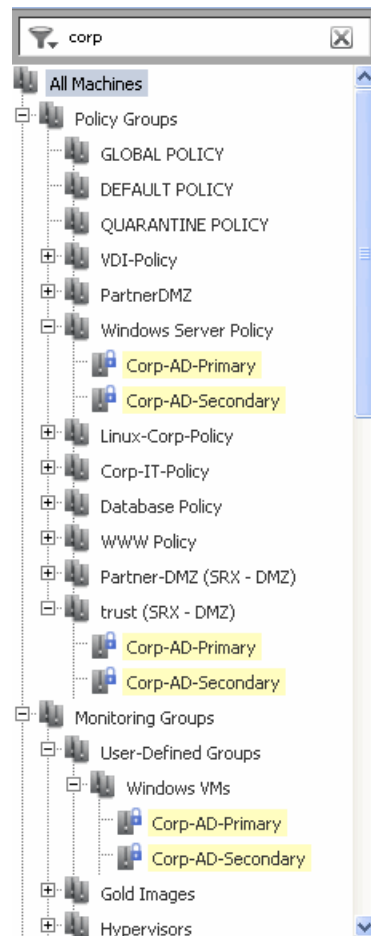
- Related Documentation**
- [About the vGW Security Design VM Tree on page 73](#)
 - [Understanding the vGW Security Design VM on page 67](#)
 - [Understanding vGW Series on page 3](#)

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 37 on page 73](#) shows VMs belonging to three groups.

Figure 37: VM Tree with Selected VMs



See “[Understanding the vGW Security Design VM Taskbar](#)” on page 71.

- [VM Tree Overview on page 73](#)
- [Locating VMs in a Complex VM Tree on page 75](#)

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 4 on page 75](#) identifies the icons that show the state of monitored VMs.

Table 4: Virtual Machine State Icons

	The VM is being fully monitored, but it is not secured. For example, no firewall policy is loaded.
	<p>The VM or the externally defined machine is not being monitored, and it has not been moved to a network secured by vGW Series.</p> <p>NOTE: Network reports can display sessions between an unmonitored system and a monitored VM.</p>
	<p>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.</p> <p>TIP: You can manually define an IP address by selecting the Settings module's vGW Application Settings > Machines.</p>
	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.



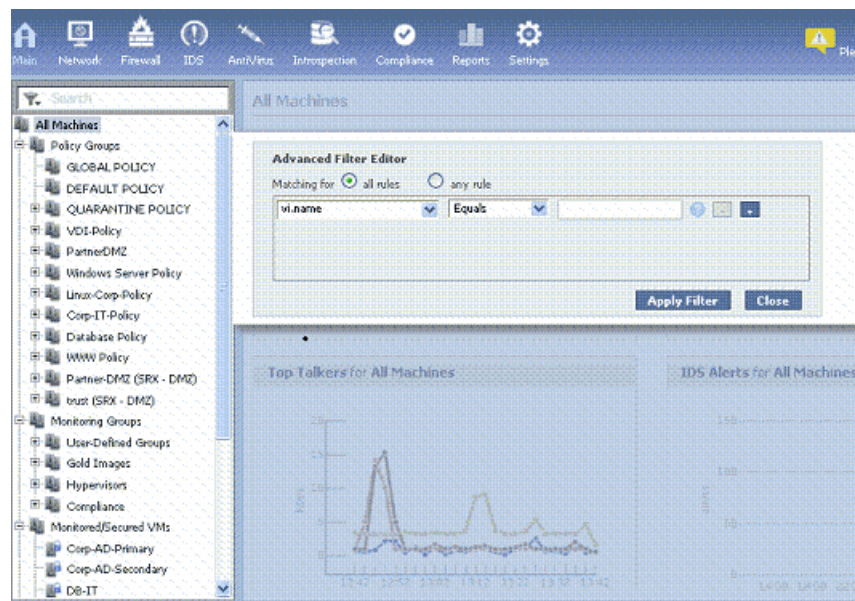
NOTE: An x icon is shown at the right side of the search field as the filter is being applied. You can use it to clear the filter.

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.

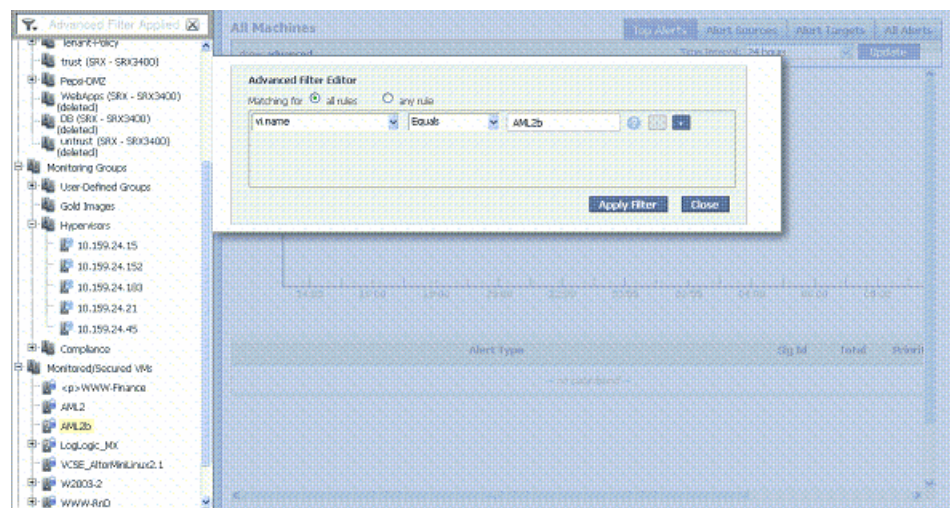
To use the advanced filter, click the icon at the left side of the search filter. This displays the Advanced Filter Editor shown in [Figure 38 on page 76](#).

Figure 38: 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 39 on page 76](#).

Figure 39: 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 71](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 6

vGW Series Main Module

This chapter describes the vGW Security Design VM Main module components. It includes the following topics:

- [Understanding the vGW Series Main Module on page 77](#)

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 77](#)
- [Status Tab on page 78](#)
- [Events and Alerts Tab on page 80](#)
- [Quarantine Tab on page 83](#)

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 40 on page 78](#).

Figure 40: 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.

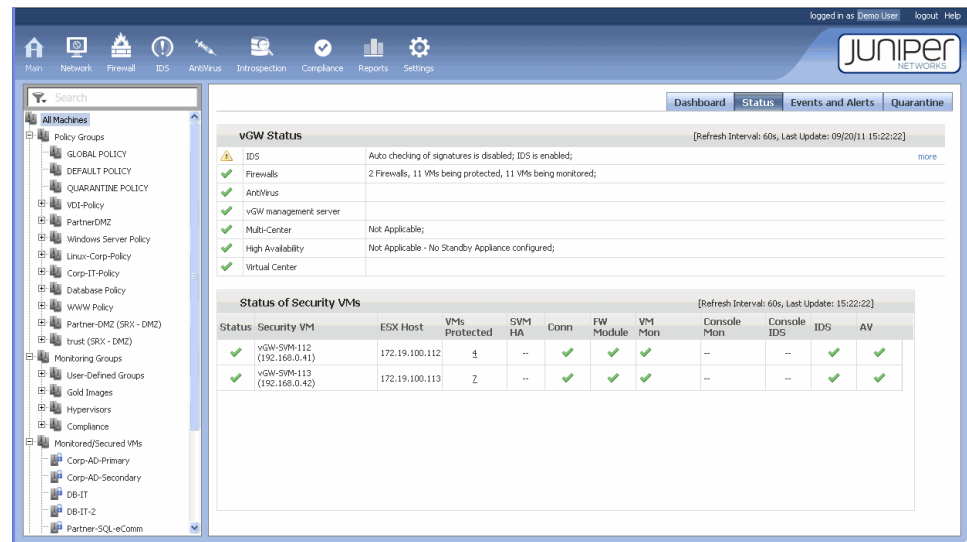
Top Talkers for All Machines—Displays network activity for the last hour.

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 41 on page 79](#).

Figure 41: 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 5 on page 79](#).

Table 5: 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 42 on page 80](#) 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 42: 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.

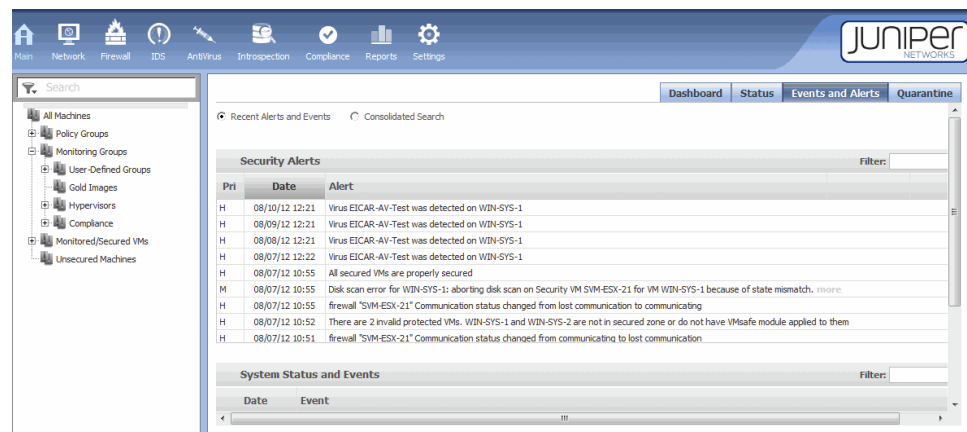
- If vGW Series HA is enabled.
- If the vGW Security VM is connected to the vGW Security Design VM.
- If the firewall module is enabled.
- If VM monitoring is used.
- IP address of the vGW Security Design VM management center.
- If IDS is used, the IP address of the IDS console.
- If IDS is enabled, 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” on page 286](#).

Events and Alerts Tab

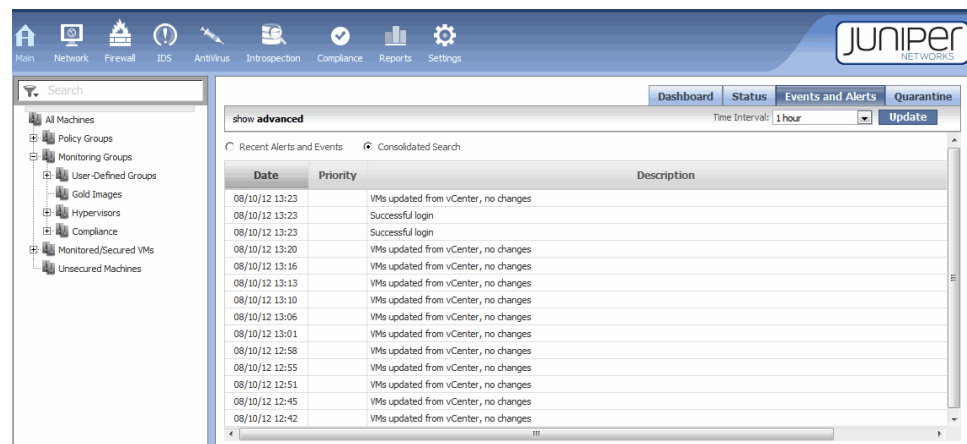
The Events and Alerts page allows you to view Security Alerts and System Status and Events messages individually, in separate panes of the page. You can use an individual filter to search each set separately. See [Figure 43 on page 81](#).

Figure 43: 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 44 on page 81](#).

Figure 44: Consolidated Logs for Events and Alerts



- [Security Alerts on page 81](#)
- [System Status and Events on page 82](#)

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.

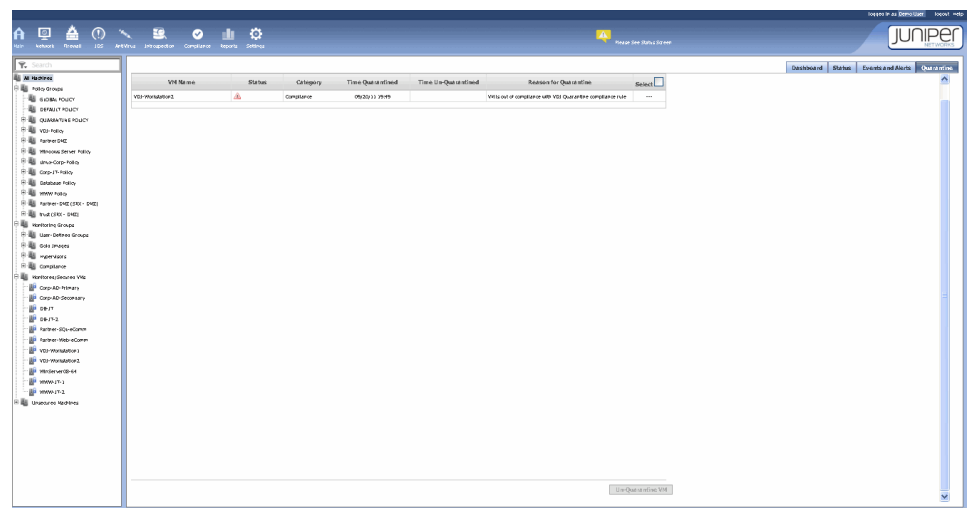
Events are listed chronologically. The events that occurred most recently are listed at the top of the table. To view additional events, you can access the vGW Security Design VM database.

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 45 on page 83](#).

Figure 45: 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” on page 131](#).

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” on page 157](#).

**Related
Documentation**

- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 7

vGW Series Network Module

This chapter covers the Network module of the vGW Security Design VM that displays information about network traffic. It includes the following topics:

- [Understanding the vGW Series Network Module on page 85](#)
- [Using the vGW Series Network and Firewall Modules Cooperatively on page 90](#)

Understanding the vGW Series Network Module

The vGW Security Design VM Network module displays network traffic for virtual machines (VMs) that are selected in the VM tree. You can view network traffic for all VMs or specific ones.

This topic includes the following sections:

- [Network Module on page 85](#)
- [Manipulating Displayed Information on page 86](#)

Network Module

The Network module contains the following six tabs:

- Summary
- Top Protocols
- Top Sources
- Top Destinations
- Top Talkers
- Connections

To display information for a VM, the VM must have a known IP address. The IP address is determined automatically if VMware Tools is installed on the VM. If it is not set automatically, you can set the IP address manually using the Settings module vGW Application Settings > Machines page.

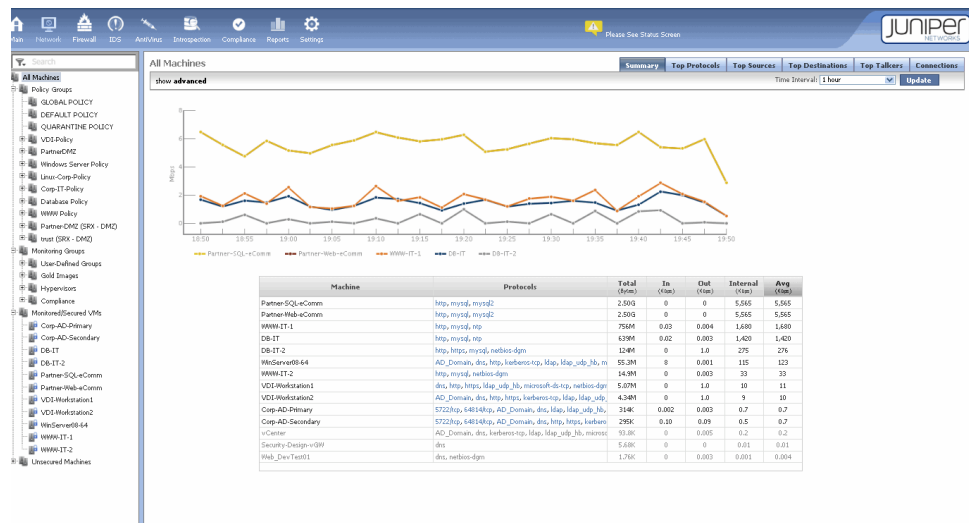
The Network module analysis takes into account IPv4 traffic and IPv6 traffic. Tables shown on the Network module tabs display information for objects with IPv4 and IPv6 addresses.

Manipulating Displayed Information

The Network Summary tab allows you to display information about all VMs, as shown in Figure 46 on page 86.

A line graph displayed at the top of the page plots bandwidth usage for the top VMs in the report. A table below the graph provides detailed network data for VMs selected in the VM tree. In this case, data for 1 hour is displayed.

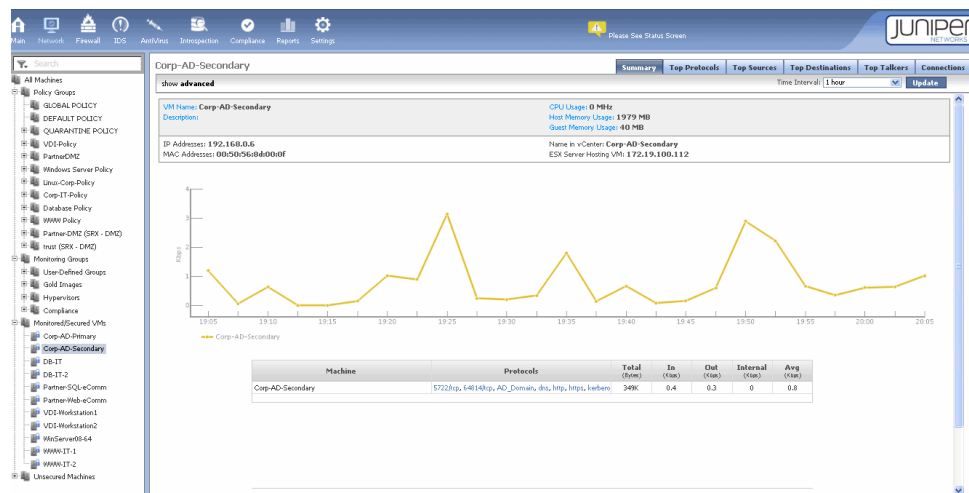
Figure 46: Network Summary Tab for All VMs



To display information about a single VM, select the VM in the VM tree.

Figure 47 on page 86 shows the information displayed for the Corp-AD-Secondary VM.

Figure 47: Main Module Network Module Summary Tab for a Single VM



To view a VM's connections, click an individual line in the graph. To display a filter for a protocol, click the protocol field.

Changing the Time Interval for Displayed Information

To change the period for which network data is plotted, use the Time Interval menu. Choose a different interval, and click **Update**. You can select a time interval or specify a custom period.



TIP: The time interval feature is also available for other vGW Security Design modules.

Figure 48 on page 87 and Figure 49 on page 87 show information for all machines for two different time periods.

Figure 48: Displaying Network Data for Different Time Intervals: Part 1

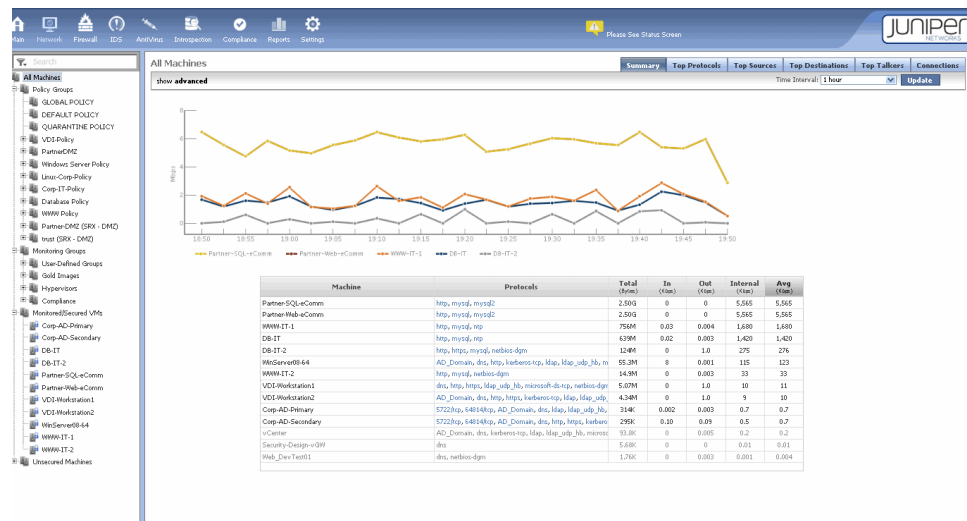
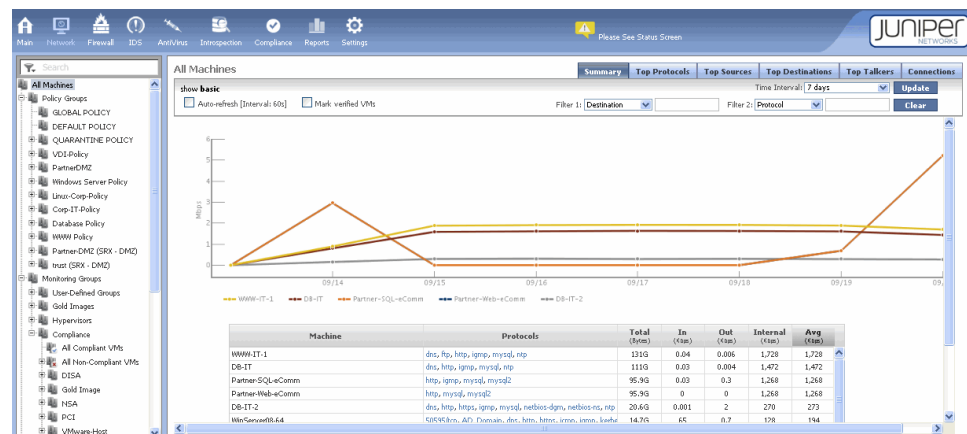


Figure 49: Displaying Network Data for Different Time Intervals: Part 2

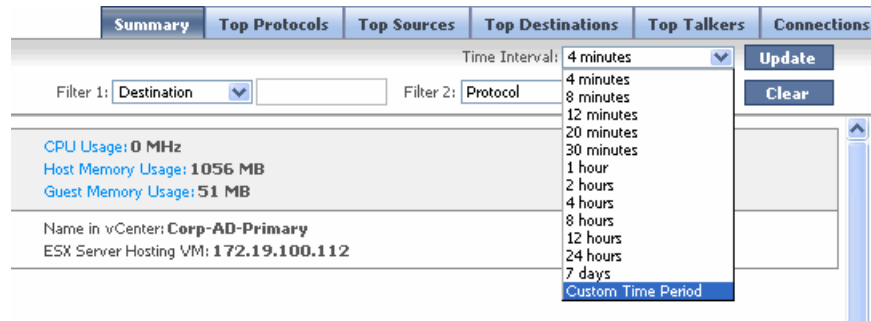


Real-time data from the last traffic interval populates the Total, In, Out, and Internal table columns. If you are charting protocols, sources, destinations, or top talkers, the interval selected is used to calculate the minimum, maximum, and average figures in the

table shown below the graph. For example, if you select 4 minutes as the time interval, the graph would show a sample of the throughput every 10 seconds. Each dot represents the average throughput value for that period.

The Custom Time Period feature allows you to view historical data. To use it, in the Time Interval menu, select **Custom Time Period**. (Figure 50 on page 88 shows the Custom Time Period menu item.)

Figure 50: Selecting a Time Interval



The custom time period is interpreted as follows:

- You can not set the custom time period to a range of less than 1 minute.


If you enter the same value for the **From** and **To** fields—that is, the same beginning and end—vGW Series automatically changes the time interval to 1 minute before the specified time.

For example, if you set the **From** field value to 01/02/13 00:00 and the **To** field value to 01/02/13 00:00, vGW Series changes the **From** time to 01/01/13 23:59 (11:59 P.M.) to allow for a time period of 1 minute. The **To** field is still interpreted as 01/02/13 00:00, the beginning of the next day.

- If you specify a valid time range, such as the **From** field set to 01/01/13 00:00 with the **To** field set to 01/02/13 00:00, vGW Series uses the time you specified.

Figure 51 on page 88 shows the Custom Time Period fields.

Figure 51: Setting the Custom Time Period



NOTE: Depending on the size of the database and the resources available to it, when you specify a custom time period, the vGW Security Design VM might take 30 minutes or more to chart the data and display it. When you want to examine a large data set, for example, data from a month or more, we recommend that you use the Reporting module.

Using Advanced Options for Filtering Network Data

You can filter the information to be displayed. To display filtering options, click **show advanced** at the left end of the time interval bar. Click the **Filter 1** and **Filter 2** menus to select filtering options and enter associated values in the related boxes. Then click **Update** to refresh the graph and data display, based on your settings. Click **Clear** to reset filter boxes.



NOTE: Configured filters affect all data in the graph and tables.

Other advanced options differ somewhat depending on the tab you are viewing. [Table 6 on page 89](#) describes the Advanced options.

Table 6: Using Advanced Options for Filtering Network Data

Select	Action
Auto-refresh	Refreshes data automatically every 60 seconds.
mark verified VMs	<p>Causes the vGW Series to automatically use the unique VMware ID/UUID as well as the IP address to validate that connections are actually coming from the identified server. vGW Series reports on both IPv4 and IPv6 addresses.</p> <p>Using both the VMware ID/UUID and the IP address protects against security threats such as IP spoofing. VMs for which this extra validation occurs can be displayed in the interface.</p>
multicast in table	<p>Includes multicast packets when monitoring. Because multicast packets are not destined for a specific host and they are seen by all machines on the network, they are included in the connection session list for all VMs.</p> <p>However, the amount of multicast traffic can be quite large, and it can obscure sessions specific to a selected VM. To remove multicast from this view, clear the multicast in table check box.</p>

To exit advanced view, click **show basic**.

Sorting Table Data

You can sort table data in the Network page by column. Drag the pointer over the column headings. When the pointer changes to the pointing hand, click the column heading to sort.

To display information for a single VM that is listed in the table, click its entry.

Related Documentation

- [Using the vGW Series Network and Firewall Modules Cooperatively on page 90](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)
- [Understanding vGW Series on page 3](#)

Using the vGW Series Network and Firewall Modules Cooperatively

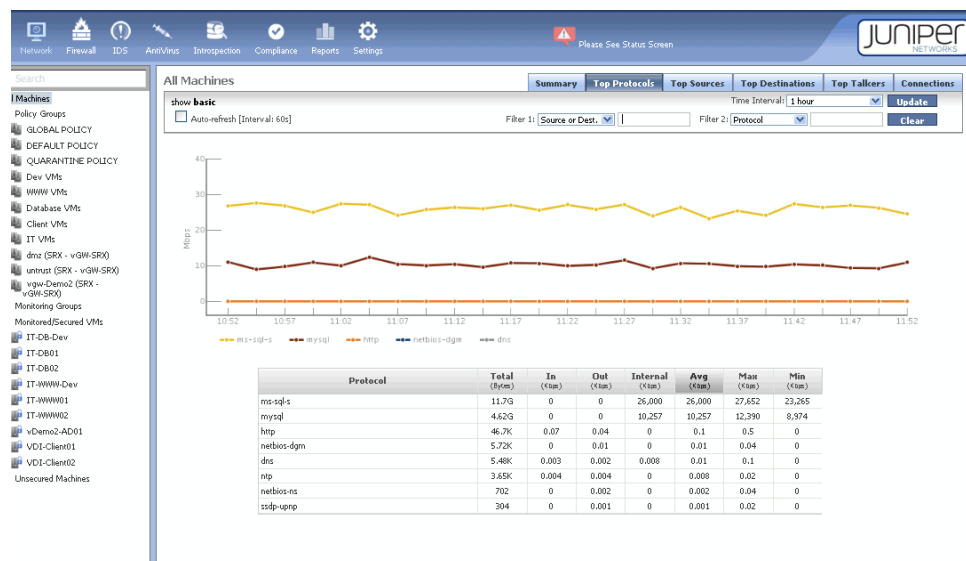
There are various ways to use the Network module in the service of the Firewall module to build a strong firewall. This topic explores some of them.

- [Network Assessment on page 90](#)
- [Using the Network Module to Observe Traffic Coming Into and Going Out from VMs on page 91](#)
- [Detecting Unexpected and Unwanted Behavior on page 91](#)
- [Using the Network and Firewall Modules Together on page 92](#)

Network Assessment

Administrators are not always aware of events that transpire on their virtualized networks because existing software for the virtualized environment does not always expose them. vGW Series Network module addresses this problem. It gives you a clear view of all traffic flows across your virtualized network. You can view overall throughput, chart protocol usage, identify sources and destinations of traffic, and identify top talkers. You can calculate minimum, maximum, and average figures across specific time intervals for these aspects of your network. In the example shown in [Figure 52 on page 90](#), the Top Protocols assessment shows that the most heavily used protocols are Microsoft SQL Server followed by MySQL. The table beneath the graph gives details on all protocols used in top down order from most used to least.

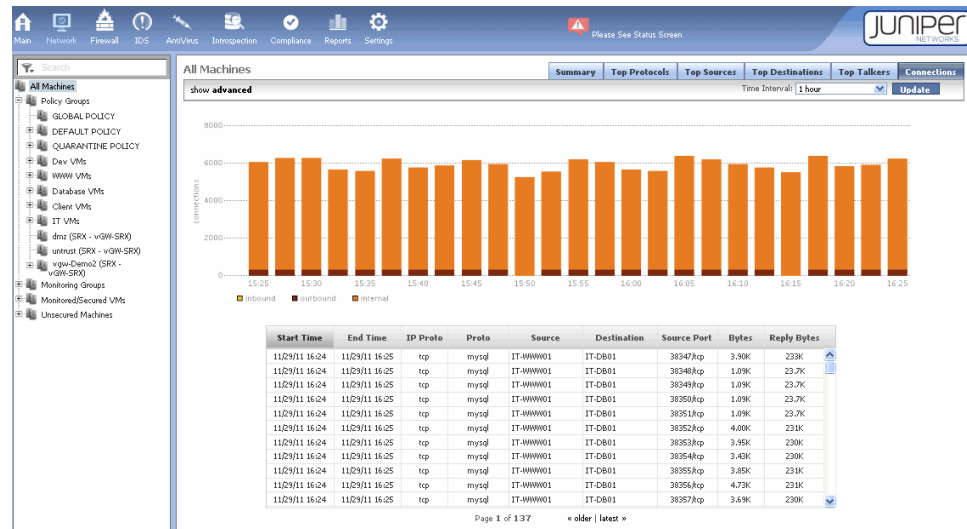
Figure 52: Top Protocols Across All Machines Example



Because the vGW Series allows you to view activity that occurs inside the hypervisor, you can quickly discover who is communicating with whom. If you were to use only the vGW Series ability to view connections in real time, you would still be able to make realistic network assessments. But the vGW Series can contribute much more information to use in your network assessment.

As [Figure 53 on page 91](#) shows, the Network module's Connections tab displays the number of connections in your network across time for all machines, whether the connections are inbound, outbound, or internal. The table beneath the graph shows when the connection was set up and when it ended, the protocol used, the source and destination endpoints, and the bytes transmitted. You can view this kind of information for an individual VM by selecting the VM in the VM tree.

Figure 53: Network Module Connection Tab Information



Using the Network Module to Observe Traffic Coming Into and Going Out from VMs

The Network module contributes to your ability to create strong firewall security in many ways. It displays information about all traffic, including traffic internal to a VM, traffic in and out of its vNICs, traffic from another VM on the same host, traffic between VMs on different hosts, and even traffic transmitted through a physical connection. In its simplest sense, you can think of this aspect of the Network module as akin to a packet sniffer, but it is far more than that.

When you use the Time Interval field to select a different time period, vGW Series redraws its graphs to let you view traffic patterns that occur during that period. You might want to use this feature to compare activity during one period of time with another, to look at past behavior, or to hone in on a VM to view its activity during a specific period.

For example, you could view all HTTP connections, the engaged workstations, and how much traffic is transmitted. You could do this for a two-day period, then a week, and then longer to observe anomalies that might exist.

Detecting Unexpected and Unwanted Behavior

The Network module can reveal unwanted behavior on your network that should be prohibited or investigated further. There are many examples of the kinds of information that the Network module might reveal. For example, you might notice that:

- Traffic might be transmitted on a particular protocol that is unusual or inappropriate, therefore raising questions.

- The protocol 999TCP might be connecting to the finance server, an unwanted event that you want the firewall to protect against.
- HTTP traffic might be transmitted to a VM that should not receive it.
- Some workstations might pull updates from a Microsoft server unintentionally instead of from local update servers.
- Thirty different protocols might be used, not all of which you were informed about. You might want to prohibit use of some of them.

Using the Network and Firewall Modules Together

When used together, the Network module and the Firewall module allow you to implement appropriate, strong security for your virtualized environment. By using the Network module to view how VMs behave in real time, you can better analyze your current security posture and observe its weaknesses.

As you begin to lock down your system through the Firewall module, the Network module becomes increasingly useful. After you use the Firewall module to refine your security policy, you can return to the Network module to determine if the change in policy produces the expected behavior.

You might still notice traffic that should not be allowed. In that case, you can return to the Firewall module, create a rule or modify an existing one, and then look at the behavioral results again in the Network module.

You can cycle through this process as many times as necessary to put in place the desired security policy. You can continue to use the Network module and the Firewall module together to implement the security you desire as your network expands and as its security requirements change.

Related Documentation

- [Understanding the vGW Series Firewall Module on page 93](#)
- [Understanding the vGW Series Network Module on page 85](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 8

vGW Series Firewall Module

This chapter covers the vGW Series Firewall module. It includes the following topics:

- [Understanding the vGW Series Firewall Module on page 93](#)
- [Understanding How vGW Series Handles ICMPv6 Protocol Traffic on page 106](#)
- [Understanding Predefined Objects for vGW Series Firewall Policy Terms on page 110](#)
- [Configuring vGW Series Firewall Policies on page 114](#)
- [Understanding vGW Series Predefined Firewall Policy for Its Components on page 120](#)

Understanding the vGW Series Firewall Module

This topic covers the vGW Series Firewall module that allows you to create reusable and individual policy rules to use in building policies for groups of VMs and individual VMs. You also use the Firewall module to apply those policies to VMs.

Before it covers the Firewall module interface, this chapter explains the policy module concepts that are fundamental to constructing firewall policies.

This topic contains the following sections:

- [The Firewall Module and the VM Tree on page 93](#)
- [Overview of the Firewall Policy Model on page 94](#)
- [Global Policy, Group Policy, and Individual VM Policy Tiers on page 95](#)
- [Firewall Policy Structure and Policy Rules Precedence on page 98](#)
- [Viewing the Complete Policy Rule Base for a VM on page 100](#)
- [The Manage Policy Tab on page 100](#)
- [The Apply Policy Tab on page 104](#)
- [The Logs Tab on page 105](#)

The Firewall Module and the VM Tree

The Firewall module of the vGW Security Design VM allows you to define, apply, and monitor security policies. To change the data displayed on a Firewall module page, select all, one, or more than one VM in the VM tree. If you select one or more VMs, but not all, information pertaining to only the selected VMs is displayed. [Figure 54 on page 94](#) shows information for a single VM.

Figure 54: Firewall Module Policy for a Single VM



Overview of the Firewall Policy Model

Security administrators of virtualized data centers invest a great deal of time and effort in planning their virtual infrastructures and building them out into group structures and categories to segment their VMs appropriately. The firewall policy model that they use to secure their virtualized infrastructure must be designed to accommodate the complexities that are intrinsic to the data center. Defining policy rules and building a firewall inside the middle of the data center differs in fundamental ways from building a perimeter firewall. Additionally, security for the virtualized data center infrastructure includes many challenges not the least of which is management of firewall policies for a large number of VMs.

The vGW Series Firewall policy used to secure the virtualized data center is modeled on the data center infrastructure overall, and it is purpose-built to meet its requirements.

- It entails group policy constructs to address group structures.
- It provides a means of simplifying the daunting task of creating policies for a large and increasing number of individual VMs.

You can create reusable policies to apply across all VMs and groups of VMs, and you can define policy rules for individual VMs.

- It allows for flexible nesting to let you define policy rule precedence within these structures as they apply to an individual VM. You can change the order of rules within global, group, or individual sets of rules to control the effect of the policy.
- It addresses the need to build flows between different systems with greater granularity than a perimeter firewall design would entail.

Ultimately every VM has its own complete firewall policy, which is composed of some or all of these parts:

- Rules that apply to all VMs in your environment. Every VM policy contains them (Global Policy rules).
- Rules that apply to the individual VM *and* others like it, if a VM belongs to a group (Group Policy rules).

- Rules that apply only to that VM, if any are required (individual VM Policy rules).

If a VM contains multiple vNICs, you can define separate policy rules for individual vNICs. These policy configurations show up in the VM rules section. See [“Configuring the vGW Series Policy per vNIC Feature” on page 261](#).

The combination of these parts gives a VM a unique firewall rule base.

Global Policy, Group Policy, and Individual VM Policy Tiers

As with many firewall designs, the vGW Series firewall policy rules are applied in a top-down fashion. To ease management of a large number of VMs and to give you control over when rules are applied, the vGW Series firewall policy allows you to define policy at three tiers: the Global Policy tier, the Group Policy tier, and the VM Policy tier. You create a Global Policy and one or more Group Policy rule sets separately. vGW Series nests them appropriately for the individual VM when you create its policy. You can move policy rules within a tier to change precedence, controlling the order in which rules are executed.

At first glance the vGW Series firewall policy nesting model might seem complex, but its simplicity and usefulness become evident as you become familiar with the symmetry at the Global Policy and Group Policy tiers and the precedence relationship within a tier and among the tiers. The Global Policy tier has high-level and low-level sections that bound the policy; the Group Policy tier is nested within the Global Policy tier and it too has high-level and low-level sections. Individual VM Policy rules are nested at the center of a VM's policy between the Group Policy high-level and low-level sections.

Although a VM policy could contain policy rules at all three tiers, it is not necessarily the case. The following sections cover each of the policy tiers in particular, but to gain an overall sense of how they can be combined to create a policy consider the following:

Ultimately every VM has its own complete firewall policy, which is composed of some or all of these parts:

- Rules that apply to all VMs in your environment. Every VM policy contains them (Global Policy rules).
- Rules that apply to the individual VM *and* others like it, if a VM belongs to a group (Group Policy rules).
- Rules that apply only to that VM, if any are required (individual VM Policy rules).

If a VM contains multiple vNICs, you can define separate policy rules for individual vNICs. These policy configurations show up in the VM rules section. See [“Configuring the vGW Series Policy per vNIC Feature” on page 261](#).

Global Policy and Group Policy rule sets contain Inbound and Outbound parts.

Global Policy

You define a reusable Global Policy whose rules apply to every VM in your environment once—it is *global*. In that it is included in every VM's policy, the Global Policy is very powerful.



NOTE: Although it is possible to delete all rules from the Global Policy, the concept of the Global Policy as applied before any other rules in the policy remains enforced. If you deleted all global rules, an empty Global Policy would be applied to the VM.

Not to diminish their usefulness, you should take care in creating rules at the Global Policy level for the very fact that they are inherited by everyone.

Both the Inbound and Outbound parts of a firewall policy contain Global Policy sections. As is the case with many firewall configurations, by default the Global policy is restrictive. It is configured to allow inbound DHCP traffic and then to reject all other inbound traffic.

You can think of the Global Policy as a template or a container for the other nested parts that will compose the entire firewall policy for any VM, keeping in mind that the Global Policy itself consists of rules.

For both the Inbound and Outbound parts of a firewall policy, the Global Policy is segmented into the following two sections:

- High-level Global Policy rules

These rules are positioned at the top of each part of a firewall policy. They are always applied to every VM first, whether that VM belongs to a group or is an individual VM. You use high-level Global Policy rules to enforce policy that cannot be overridden by any individual VM Policy rule.

For example, in addition to enforcing corporate policy, you might use high-level Global Policy rules to prevent outbreaks and protect against vulnerabilities. You might add a Global Policy rule to block access to a vulnerable service until it is updated with all of the required patches.

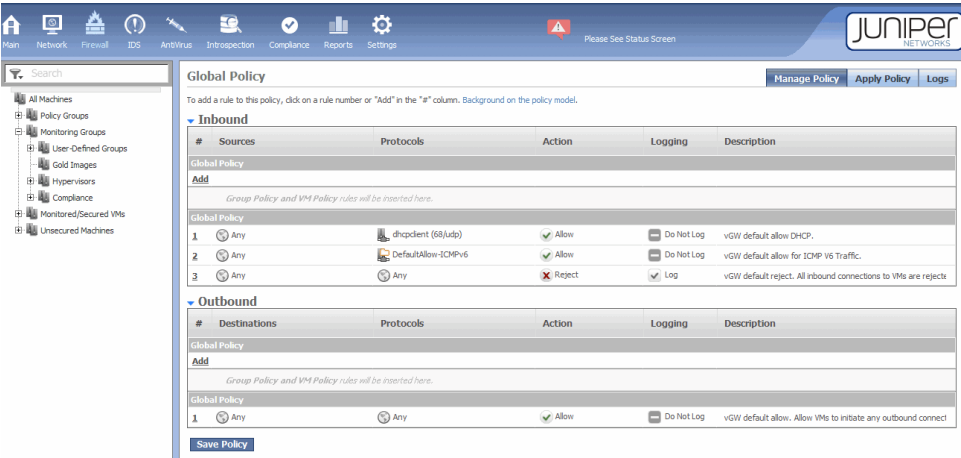
- Low-level Global Policy rules

These rules are positioned at the bottom of each part of a firewall policy. In any overall individual VM's firewall policy, they are applied last. They are applied to every VM. For example, for the Inbound part of a Global Policy, if an incoming connection is processed according to the appropriate firewall policy and it does not match any of the preceding rules, it falls through to the Inbound low-level Global Policy rules. Low-level Global policy rules are typically used as clean-up rules. By default, the Inbound low-level Global Policy rule rejects all connection attempts. It is defined as any-any-reject.

Between the high-level and low-level sets of Global Policy rules is a placeholder that allows for nesting of Group Policy rule sets and individual VM Policy rules.

To create a Global Policy, you select **GLOBAL POLICY** under Policy Groups in the VM tree. The page shown in [Figure 55 on page 97](#) is displayed.

Figure 55: Global Policy



Group Policy

Most of the daily policy management that security administrators of virtualized environments carry out is at the group level. Most likely you have structured your environment along lines of groups of VM with similar characteristics and you want to apply a similar policy to VMs that are members of a group.



NOTE: In the nested model, a VM might belong to a Policy Group and inherit the Group Policy rules defined for that group, but it also might have its own individual VM Policy rules that contribute to its overall firewall policy rule base.

For example, you might organize VMs into functional groups such as Web servers and database servers, and you might want to apply a different set of policy rules to each group. In your environment, you might create different groups for MS Windows systems versus Linux systems. To apply the appropriate security, you could define a different Group Policy for each of them.

The Group Policy concept allows you to define policy rules that are relevant to the VMs that comprise the group. As new VMs are created and added to a Policy Group, the Group Policy associated with the group is applied to them.

A VM might belong to multiple Policy Groups. For example, a VM might be a Windows VM and belong to the Windows group, but it also might be used as a Web server and belong to the Web servers group. In this case, the VM gets the Group Policy rules for both groups.

Individual VM Policy Rules

At the center of the entire firewall policy for an individual VM are any particular VM Policy rules that you define for that VM. Until this point, the firewall policy for an individual VM is composed of reusable parts—the Global Policy and, if the VM belongs to any Policy Groups, Group Policy rules.

You can apply individual VM Policy rules to a VM policy for particular purposes that distinguish that VM's policy from others. For example, you might want RADIUS access to a VM that is not applied at the Global Policy or Group Policy levels. To accomplish that, in the VM's firewall policy, you would define an Inbound VM Policy rule that allowed RADIUS access to the VM.

Default Policy

A newly created VM that does not have individual policy rules or group policy rules associated with it is automatically assigned the Default Policy. Also, when the policy for a VM includes one or more VM Policy rules but it does not include Group Policy rules, the VM inherits the Default Policy rules, in addition to the individual ones. Later if it becomes a member of a group, then it inherits that group's Group Policy rules, and the Default Policy rules no longer apply.

By default, the Default Policy does not contain any policy rules. It is assumed that you will define the policy that you want to use as the default.

Quarantine Policy

When a VM is infected by a virus and the scanning configuration specifies "Quarantine the VM", the VM is put in the Quarantine policy group. The Quarantine Policy that you define is applied to all VMs in the Quarantine policy group. When you remove the VM from the group, the Quarantine policy is removed.

To remove the VM from the Quarantine policy group, use the Main module Quarantine tab. Select the VM, and click **Un-quarantine**.

For details on how the parts of the quarantine process work together for a quarantined VM, see "Understanding Quarantined VMs and How to Manage Them" on page 152.

Firewall Policy Structure and Policy Rules Precedence

The vGW Series Firewall policy model is premised on a pre-post concept that allows you to manage rules execution precedence.

Consider the nested structure of a firewall policy. To summarize the order, a firewall policy has inbound and outbound sections. The Inbound section contains the high-level Global Policy rules followed by, the Group Policy rules, then the individual VM Policy rules, and finally the default Global Policy rules. The default Global Policy rules consist of a rule to allow DHCP traffic, a rule to allow certain types of ICPMv6 traffic, and, at the bottom, a rule to reject all other inbound traffic. The outbound section contains the same parts in the same order, only its Global Policy section contains a single rule that allows VMs to initiate outbound connections.

high-level Global Policy— At the top of the Inbound section is the high-level Global Policy tier, containing any global policies that you add.

high-level Group Policy—Beneath it is the high-level Group Policy section containing any of Policy Groups rule sets that apply to the individual VM that you want executed *before* the individual VM Policy rules.

VM Policy—Beneath it is the high-level VM Policy section containing any individual rules that you define for the VM whose policy you are creating.

low-level Group Policy—Beneath it is the low-level Group Policy section containing any group rule sets for the VM that you want to be executed *after* its individual ones.

Default Global Policy—The default Global Policy rules consist of a rule to allow DHCP traffic, a rule to allow certain types of ICPMv6 traffic, and, at the bottom, a rule to reject all other inbound traffic.

It is this structure that allows you to manipulate the order in which rules are executed for the individual VM firewall policy. The vGW Series Policy model affords you extensive, flexible control over the order in which rules are executed. You can move rules up and down within their sets; you can move rules from a low-level section of one tier to that tier's high-level section or the opposite, and you can reorganize individual VM Policy rules.

Rules are executed in a top-down fashion:

- High-level Global Policy rules are always executed first, and that cannot be changed. However, you can manage the order in which Global Policy rules are executed by moving them up and down in the set.
- High-level Group Policy rules are executed next. They are always executed before individual VM Policy rules, but you can also change the order in which they are executed by moving them up and down within the set.
- Individual VM Policy rules are executed next, and you can change their order to control when they are executed.
- Low-level Group Policy rules are always executed after the individual VM Policy rules.

By placing some of the Group Policy's rules in its low-level section, you are able to specify that in most cases you want these rules applied to all VMs that belong to the Policy Group *after* the individual VM Policy rules are executed. You will allow VM Policy rules for individual VMs to take precedence over these Group Policy rules.

- Finally, low-level Global Policy rules are executed for every VM.

For example:

- If you move a rule *up* from its low-level Group Policy section to its high-level counterpart, that rule is executed *before* any individual VM Policy rule, and it *cannot* be overridden by a VM Policy rule. Previously, when it resided in the low-level Group Policy section, a VM Policy rule could override it.
- If you move a rule *down* from its high-level Group Policy section to its low-level counterpart, that rule is executed *after* any individual VM Policy rule, and it *can* be overridden by a VM Policy rule. Previously, when it resided in the high-level Group Policy section, a VM Policy rule could not override it.

When you nest rules for a VM's firewall policy, take into account precedence among the various levels of the policy. For example, consider a policy for a VM whose inbound low-level Group Policy section includes a rule that allows management access to the

VM. Suppose that as the data center administrator you will always want management access to the VM. However, you understand that another administrator could create a firewall policy intended for an individual VM that is a member of the Windows VMs group as part of the group policy. That administrator could define a VM Policy rule for the individual VM that would reject management access to the VM, effectively denying you access. Because the Group Policy rule allowing access is in the low-level section of the Group Policy rule set, the individual VM Policy rule would override it.

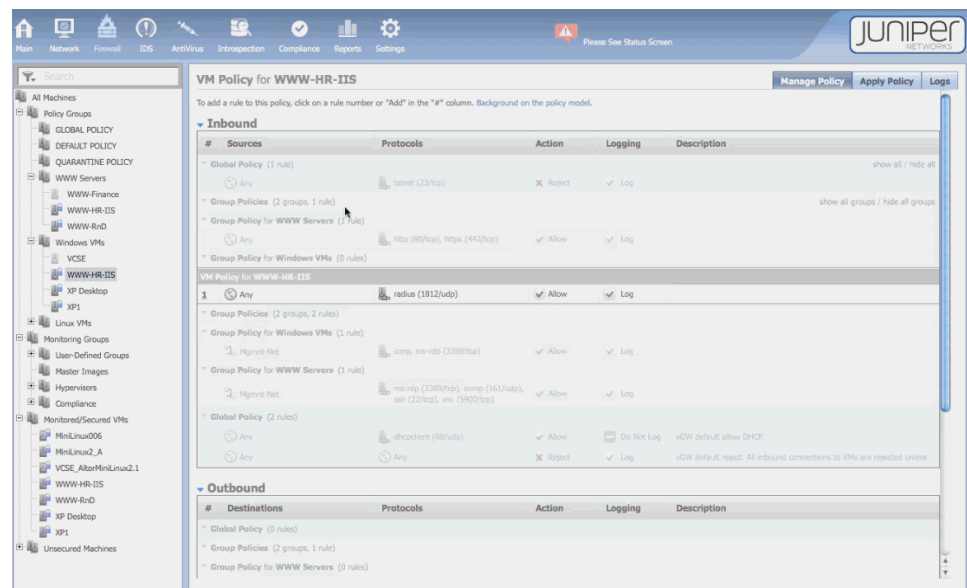
To ensure that you always have management access, you could affect the precedence in the policy for any VM that belongs to that group by moving the rule that allows management access up from the low-level Group Policy section to the high-level Group Policy section. To do so, click the rule number in the low-level Group Policy and select **Move Rule Up** from the list.

Viewing the Complete Policy Rule Base for a VM

Each VM protected by a vGW firewall policy can be thought of as having its own firewall policy. The resulting full policy for a VM always includes a Global Policy, Group Policies if the VM belongs to Policy Groups, and individual VM Policy rules that are specific to it.

After you have created a firewall policy for a VM or you want to understand its policy, you can expand it to see its entire rule base. To do this, select the Firewall module. In the VM tree, select the VM. On the upper-right side of the VM Policy page, click **show-all**. See [Figure 56 on page 100](#).

Figure 56: VM Policy Expanded Rule Base

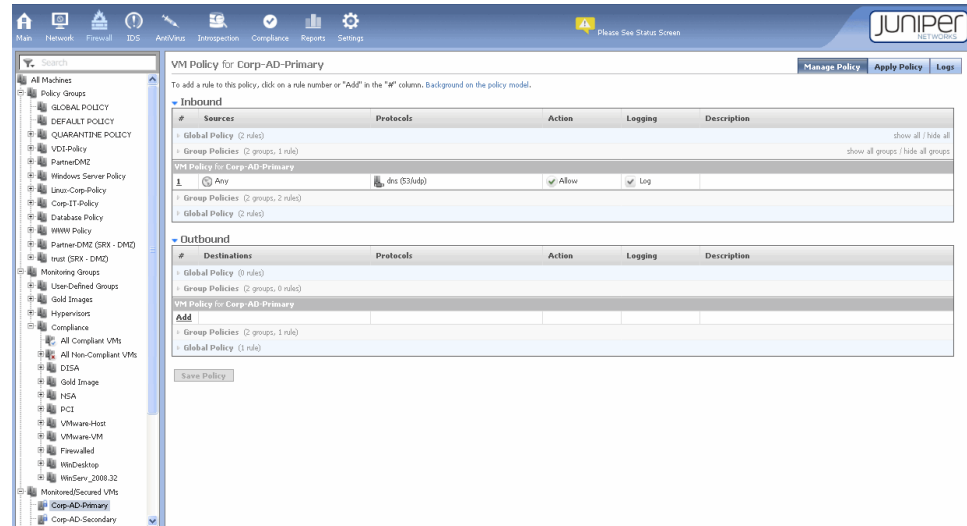


The Manage Policy Tab

The Manage Policy tab allows you to define and edit security policies. The Manage Policy page shows the policy configured for the group of VMs or the VM that is selected in the VM tree. To change the data displayed on the Manage Policy page, select a different

object in the VM tree. You can select all machines, a group, or an individual VM. [Figure 57 on page 101](#) shows the policy for the Corp-AD-Primary VM.

Figure 57: Firewall Module Manage Policy Page



This section contains the following parts:

- [Policy Per vNIC and Dual Stack on page 101](#)
- [Creating a Policy Rule on page 101](#)

Policy Per vNIC and Dual Stack

A single VM may have multiple vNICs attached to it. In the case of a dual stack, a VM would have a vNIC with an IPv4 address and an IPv6 address bound to it.

vGW Series provides a feature called Policy per vNIC that allows you to define separate policies for individual vNICs attached to the same VM. You can configure separate policies for individual vNICs, separate policies for some of them while leaving others unsecured, or you can use the same policy for all of them.


Using the Policy per vNIC feature, you can handily apply different policy rules to vNICs passing IPv4 traffic from those used for IPv6 traffic even when the vNICs are attached to the same VM. To apply the rule to all traffic of a type, you could use the predefined terms **Any-IPv4** and **Any-IPv6**.

Creating a Policy Rule

To create a policy rule:

1. Click a rule number in the rule numbers (#) column.
2. Select **Add Rule Above** or **Add Rule Below**. See [Figure 58 on page 102](#).

Figure 58: Adding a Rule




NOTE: Rules are applied in order of execution from top to bottom.

- Configure policy settings by clicking the table cells and editing the information using the dialog box.

For example, to specify a protocol for the rule, click the default value **Any**, which displays a dialog box. To quickly make selections, type the first letter of the item that you want to select in the filter field. See Figure 59 on page 102.

Typing the letter **t** in the All Protocols dialog box scrolls to the telnet selection in the list.

Figure 59: Using the Dialog Box Filter to Add Terms for policy rules



To immediately select an item, type directly into the Filter box.

To define a policy that contains all protocols except for a few:

1. Click **Advanced** at the bottom of the dialog box.
2. Click **Negate this selection**.
As a result, “All protocols except” is displayed at the top of the Selected Protocols list.
3. For each protocol or protocol group that you want to exclude from the policy rule, select the object and click the right arrow to move it to the list.
4. Click **Apply**, when you are finished.
5. When you have finished entering or editing all policy settings, click **Save** to save your changes in the vGW Security Design VM database.



WARNING: For new policy rules to take effect, you must apply the policy changes using the Apply Policy tab. You can apply rules immediately or during maintenance.

To delete or disable/deactivate an existing rule, click the rule number and choose the appropriate option. Disabled rules appear dimmed and are shown with a strike-through mark.

[Table 7 on page 103](#) describes the policy configuration settings.

Table 7: Firewall Policy Configuration Settings

Field	Function
Sources	Define the object from which the connection originates.
Protocols	Define which protocols are used in the rule. You can also dynamically create a new protocol or protocol group by selecting the appropriate option.
Action	Allow the connection, drop the connection (silent drop), or reject the connection (drop traffic and send source a notification). In addition, you can redirect or duplicate packets to third-party devices using Settings > Security Settings > Global > External Inspection Devices. See “Configuring Global Settings Using the vGW Series Settings Module (VMware)” on page 282 .
Logging	Log the connection matching the rule, skip logging for this connection, or send an alert when this connection matches the rule. The Alert option directs the vGW Series to send e-mail messages or SNMP traps. See “Alerts” on page 80 .
Description	Enter a description for the policy.

The Apply Policy Tab

The Apply Policy tab allows you to push security policies out to the vGW Security VM firewall to protect the VMs in your infrastructure. When you create or modify a policy, it is not applied to the VM automatically. For new policy rules to take effect, you must apply the policy changes using the Apply Policy tab. You can apply rules immediately or during maintenance.

You use the VM tree on the left side of the Apply Policy page to select the VMs to apply policies to.

Reflecting the hierarchy in which you create a VM policy, the Apply Policy table shows:

- That the VM has a Global Policy, its Group Policies, if it belongs to a group, and any individual policies configured specifically for it.



NOTE: If there are no Group or individual policy rules for a VM, the Global Policy is applied.

- If a VM has multiple vNICs, whether Policy per vNIC is applied to it.
- The vGW Security VM that protects the VM.
- The date that the policy was installed.

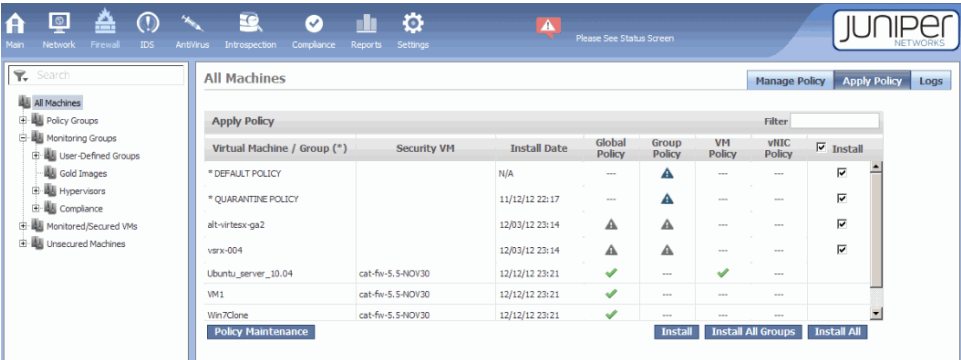
To install a policy on one or more selected VMs:

1. Select the **Install** check box at the right of the title bar.
2. Select the check box in the Install column at the right of the VM's row.
3. Click **Install** at the bottom of the page.

To install policies for all VMs, click the **Install** check box at the top of the column, then click **Install All**. To install policies for all Groups, click **Install All Groups**.






Figure 60 on page 104 shows the Apply Policy page.

Figure 60: Firewall Apply Policy Page



See [Table 8 on page 105](#) for a list of icons displayed for VMs on the Apply Policy page.

Table 8: Firewall Policy Icons

Icon	Indicates that
	The policy is current and no further actions are required.
	The VM is in a policy group, but it cannot retrieve policies because it is not protected by a vGW Security VM firewall. This usually indicates an error condition that you should investigate.
	<p>The policy type does not exist for the VM. For example, an individual VM policy for that VM is not configured.</p> <p>You are not required to build individual VM policies for each VM.</p>
	The policy has been modified, and it needs to be deployed for the VM.
	An error condition exists that prevents installation of the policy. When a policy distribution problem exists but the old policy works properly, a check mark icon might be displayed.



TIP: Place the pointer over a policy status icon to display a tool tip that describes the icon.

When you are ready to implement a policy, click either **install** or **install all** to push the policy out to the firewall. This action causes the policy to be deployed on the selected VMs or the vNICs of the VMs, if the Policy per vNIC feature is used.



NOTE: When you attempt to apply a policy to a vNIC that is not secured and that belongs to a protected VM, the policy is not applied. The following message is displayed:

“Policy was compiled and saved. This VM is currently not associated with a firewall, so the policy is not being immediately loaded on a firewall. This could be because the VMs migrated to an unprotected host or are powered off. Once the VM will be associated to a firewall, the corresponding saved policy will be enforced.”

The Logs Tab

You can define policy rules to specify Log, Don't Log, and Alert notification options. When you select **Log** or **Alert** for a rule, traffic that matches that rule is logged. [Figure 61 on page 106](#) shows the Logs tab.

For the Logs tab, you can use an advanced option that includes a mark verified VMs setting. vGW Series uses the unique VMware ID/UUID in addition to an IP address to validate that connections are coming from the identified server. This feature protects the network from issues such as IP spoofing and DHCP changes. VMs for which this extra validation is allowed are flagged with an asterisk (*). You can use the mark verified VMs setting to display or hide the icon. Click **Auto-refresh** to refresh the log displayed automatically every 60 seconds.

The log entries show both IPv4 and IPv6 addresses.

Figure 61: Firewall Module Logs Tab

Start Time	Rule Id	Action	Source	Source Port	Destination	Proto	IP Proto	Record Id
09/02/11 10:47	29	Allow	Partner-SQL-eConn	46990/tcp	Partner-Web-eConn	mysql	tcp	19399995
09/02/11 10:47	5	Reject	VD1-Workstation2	65100/tcp	WWW-TT-1	http	tcp	19399996
09/02/11 10:47	5	Reject	VD1-Workstation2	65100/tcp	WWW-TT-1	http	tcp	19399996
09/02/11 10:47	5	Reject	VD1-Workstation2	65099/tcp	WWW-TT-1	http	tcp	19399995
09/02/11 10:47	5	Reject	VD1-Workstation2	65099/tcp	WWW-TT-1	http	tcp	19399995
09/02/11 10:47	5	Reject	VD1-Workstation2	65098/tcp	WWW-TT-1	http	tcp	19399991
09/02/11 10:47	5	Reject	VD1-Workstation2	65098/tcp	WWW-TT-1	http	tcp	19399995
09/02/11 10:47	2	Allow	Partner-Web-eConn	54091/tcp	Partner-SQL-eConn	mysql	tcp	19399985
09/02/11 10:47	2	Allow	Partner-Web-eConn	54090/tcp	Partner-SQL-eConn	mysql	tcp	19399985
09/02/11 10:47	2	Allow	Partner-Web-eConn	54089/tcp	Partner-SQL-eConn	mysql	tcp	19399987
09/02/11 10:47	2	Allow	Partner-Web-eConn	54088/tcp	Partner-SQL-eConn	mysql	tcp	19399986
09/02/11 10:47	29	Allow	Partner-SQL-eConn	43672/tcp	Partner-Web-eConn	http	tcp	19399988
09/02/11 10:47	29	Allow	Partner-SQL-eConn	43671/tcp	Partner-Web-eConn	http	tcp	19399984
09/02/11 10:47	29	Allow	Partner-SQL-eConn	43670/tcp	Partner-Web-eConn	http	tcp	19399983
09/02/11 10:47	29	Allow	Partner-SQL-eConn	43669/tcp	Partner-Web-eConn	http	tcp	19399982
09/02/11 10:47	29	Allow	Partner-SQL-eConn	43668/tcp	Partner-Web-eConn	http	tcp	19399981
09/02/11 10:47	29	Allow	Partner-SQL-eConn	46304/tcp	Partner-Web-eConn	mysql	tcp	19399980
09/02/11 10:47	29	Allow	Partner-SQL-eConn	46303/tcp	Partner-Web-eConn	mysql	tcp	19399975
09/02/11 10:47	29	Allow	Partner-SQL-eConn	46302/tcp	Partner-Web-eConn	mysql	tcp	19399975
09/02/11 10:47	29	Allow	Partner-SQL-eConn	46301/tcp	Partner-Web-eConn	mysql	tcp	19399977
09/02/11 10:47	29	Allow	Partner-Web-eConn	54091/tcp	Partner-SQL-eConn	mysql	tcp	19399976
09/02/11 10:47	29	Allow	Partner-Web-eConn	54090/tcp	Partner-SQL-eConn	mysql	tcp	19399975
09/02/11 10:47	29	Allow	Partner-Web-eConn	54089/tcp	Partner-SQL-eConn	mysql	tcp	19399974
09/02/11 10:47	29	Allow	Partner-Web-eConn	54088/tcp	Partner-SQL-eConn	mysql	tcp	19399974
09/02/11 10:47	33	Allow	Partner-SQL-eConn	43672/tcp	Partner-Web-eConn	http	tcp	19399972
09/02/11 10:47	33	Allow	Partner-SQL-eConn	43671/tcp	Partner-Web-eConn	http	tcp	19399971
09/02/11 10:47	33	Allow	Partner-SQL-eConn	43670/tcp	Partner-Web-eConn	http	tcp	19399970

You can use filters to refine the display of log entries. To display only those logs related to a specific VM, select the VM in the VM tree pane.

Related Documentation

- [Understanding the vGW Series Policy per vNIC Feature on page 258](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security Design VM Taskbar on page 71](#)
- [About the vGW Security Design VM Tree on page 73](#)
- [Understanding the vGW Series Network Module on page 85](#)
- [Understanding vGW Series on page 3](#)

Understanding How vGW Series Handles ICMPv6 Protocol Traffic

This topic covers the Internet Control Message Protocol version 6 (ICMPv6) which is integral to IPv6 and fundamental to the proper functioning of IPv6 networks.

It describes the vGW Series default firewall policy protocol group for handling ICMPv6 traffic.



WARNING: By default vGW Series allows inbound and outbound ICMPv6 traffic. Juniper Networks strongly recommends that you not override this default policy because of the important role that ICMPv6 plays in establishing and maintaining communication in IPv6 networks.

- [About ICMPv6 on page 107](#)
- [Filtering ICMPv6 Packets on page 107](#)
- [Default Policy Group for Allowing Inbound ICMPv6 Packets on page 108](#)

About ICMPv6

ICMPv6 consists of a large number of messages with diverse functions which, like ICMP messages for IPv4 networks, could be categorized broadly as error and information messages.

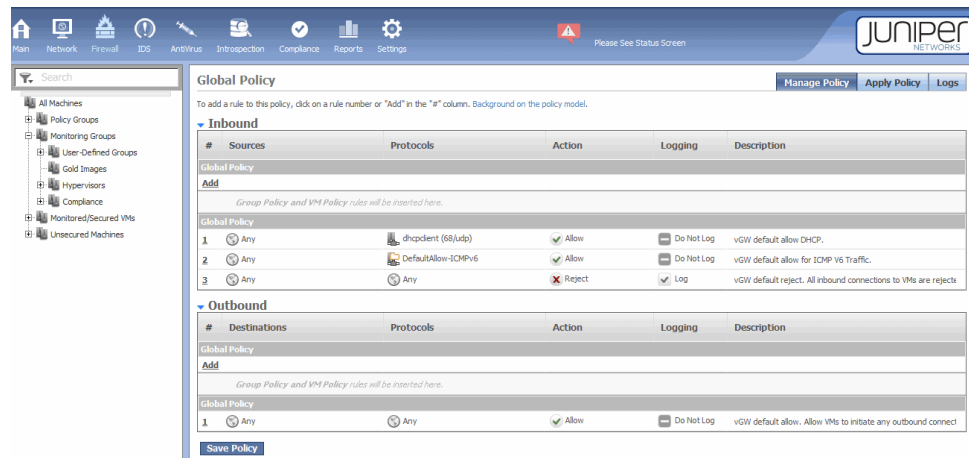
ICMP for IPv4 is an auxiliary protocol not necessarily required for IPv4 proper functioning. By contrast, ICMPv6 is an essential component in the establishment and maintenance of IPv6 communications. Among the messages it includes are those for address assignment, address resolution, and multicast group management. ICMPv6 error messages and information messages are transported by IPv6 packets in which the IPv6 Next Header value for ICMPv6 is set to 58.

Filtering ICMPv6 Packets

In IPv4 networks, it is common practice for firewalls to drop ICMP Echo Request messages to protect against scanning attacks and to minimize the risk of denial of service attacks. Port scanning in IPv6 networks is less severe, so it is not necessary to filter IPv6 Echo Requests. In practice, it is important to avoid aggressive filtering of ICMPv6 packets. Because they are fundamental to the proper functioning of IPv6 networks and tunneling, it is essential that ICMPv6 connectivity messages are allowed to pass through the firewall.

vGW Series establishes a default protocol group called DefaultAllow-ICMPv6 that allows access to traffic from a comprehensive set of ICMPv6 protocols. A default rule for the DefaultAllow-ICMPv6 protocol is created that is applied to the inbound Global policy rule set to allow this inbound traffic. See [Figure 62 on page 108](#).

Figure 62: Default Global Policy Showing Default ICMPv6 Allow Group



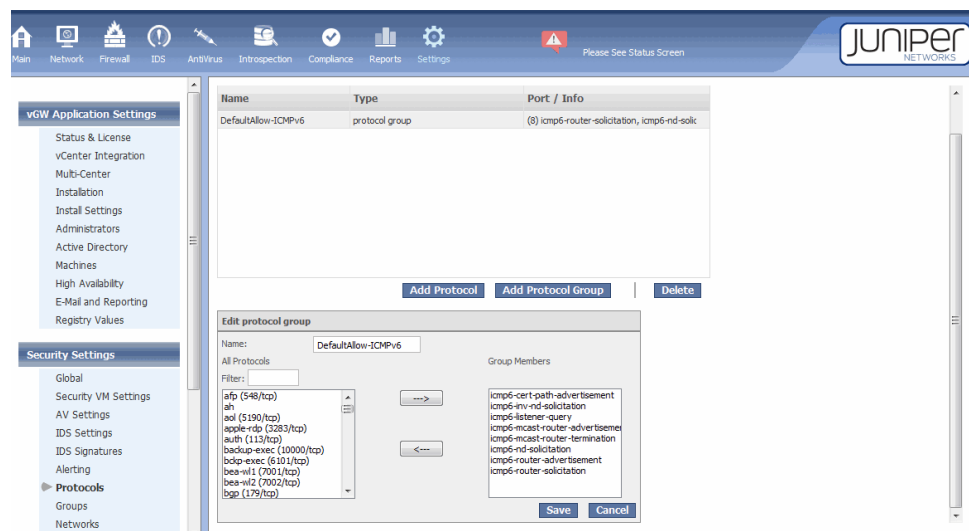
Default Policy Group for Allowing Inbound ICMPv6 Packets

vGW Series provides the predefined DefaultAllow-ICMPv6 protocol group that allows inbound ICMPv6 traffic for all types of packets included in the group. Because ICMPv6 is critical to proper IPv6 functioning, it is important that you allow this traffic. However, if for some reason you wish to block traffic from one or more ICMPv6 protocols that are members of the default protocol group, you can edit the list to exclude them from the *allow* condition and filter the traffic. See [“Editing the Default ICMPv6 Protocols Group Members”](#) on page 109.

Viewing the Default ICMPv6 Protocols Group Members

You can view the list of ICMPv6 protocols that comprise the DefaultAllow-ICMPv6 protocol group on the Settings module Security Settings > Protocols page. See [Figure 63 on page 108](#).

Figure 63: Protocols Settings ICMPv6 Default Protocol Group



To view the list:

1. Beside **Protocols**, select **Groups**.
2. Click **DefaultAllow-ICMPv6**.

The column on the right side of the Edit protocol group pane shows the group members:

- icmp6-listener-query
130. Multicast Listener Query (RFC 2710)
- icmp6-router-solicitation
133. Router Solicitation (RFC 4861)
- icmp6-router-advertisement
134. Router Advertisement (RFC 2461)
- icmp6-nd-solicitation
135. Neighbor Discovery Solicitation (RFC 4861)
- icmp6-inv-nd-solicitation
141. Inverse Neighbor Discovery Solicitation Message (RFC 3122)
- icmp6-cert-path-advertisement
149. Certification Path Advertisement Message (RFC 3971)
- icmp6-mcast-router-advertisement
151. Multicast Router Advertisement (RFC 4286)
- icmp6-mcast-router-termination
153. Multicast Router Termination (RFC 4286)

Editing the Default ICMPv6 Protocols Group Members

If you must block traffic on any of the ICMPv6 protocols in the vGW DefaultAllow-ICMPv6 protocol group, you can edit the group from Settings module Security Settings > Protocol page.

To edit the list from the Settings module Security Settings > Protocol page:

1. Beside **Protocols**, select **Groups**.
2. Click **DefaultAllow-ICMPv6**.

The column on the right side of the Edit protocol group pane shows the group members:

- icmp6-cert-path-advertisement
- icmp6-inv-nd-solicitation
- icmp6-listener-query
- icmp6-mcast-router-advertisement
- icmp6-mcast-router-termination

- icmp6-nd-solicitation
 - icmp6-router-advertisement
 - icmp6-router-solicitation
3. Select the ICMPv6 protocol that you want to remove from the list, thereby blocking its packets, and click the left facing arrow.
- Repeat this process for each protocol that you want to remove from the list.
4. Click **Save**.

Related Documentation

- [Understanding the vGW Series Firewall Module on page 93](#)
- [Understanding vGW Series Predefined Firewall Policy for Its Components on page 120](#)
- [Understanding vGW Series IPv6 Support on page 54](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Understanding Predefined Objects for vGW Series Firewall Policy Terms

This topic focuses primarily on the vGW Series predefined objects that you can use for source and destination terms in firewall policy rules. It summarizes the various ways in which you can specify addresses for these terms.

- [Defining and Selecting Source and Destination Terms for Policy Rules on page 110](#)
- [Predefined Global IP Address Objects on page 111](#)
- [Predefined Network Objects on page 111](#)

Defining and Selecting Source and Destination Terms for Policy Rules

To create firewall policies, you specify rules. You add inbound and outbound rules to a policy to specify the source and destination of traffic. You select a value for the source or the destination of a term from the list of existing objects that is displayed when you right-click the rule numbers column in the Inbound (Sources) and Outbound (Destinations) parts of a policy.

vGW Series provides the following ways in which you can define the addresses for a rule's source or destination terms:

- You can define these addresses dynamically as you create the rule. You can create groups or machines and then use them in the rule.

As a convenience, the vGW Security Design VM makes the configuration panes that you use for this purpose available from the Manage Policy page of the Firewall module that you use to define the policy. They are the same panes that you use to create the objects from other parts of the vGW Security Design VM.

- You can select a network or a machine that you have already defined.

- You can select any of the predefined objects that vGW Series provides. The following sections cover these objects.

Predefined Global IP Address Objects

vGW Series Release 5.5 introduces support for IPv6, including configuration of policies on IPv6 traffic. vGW Series provides the following predefined objects that allow you to refer to IP addresses collectively by type—whether IPv4 addresses or IPv6 addresses—in a policy rule’s source and destination terms:

Any—Matches any IPv4 and IPv6 address.

Any-IPv4—Matches any IPv4 address.

Any-IPv6—Matches any IPv6 address.

In releases earlier than version 5.5—releases before vGW Series supported IPv6—the term Any referred to any IPv4 address. For environments in which not all vGW Series components are at version 5.5 or later, the term Any also refers to any IPv4 address. It reverts back to the meaning it had in environments that support only IPv4 traffic. For more information about how Any is interpreted in mixed vGW Series components environments, see [“IPv6 Support in Homogeneous and Heterogeneous vGW Series Environments” on page 60](#).



WARNING: All vGW Series components must be at version 5.5 or later for you to be able to create policies on IPv6 traffic.

Predefined Network Objects

vGW Series provides predefined network objects for well-known IP address ranges and prefixes that you can use in policy rule terms for either source or destination addresses. It also provides network objects for other IPv6 and IPv4 addresses. This section covers both groups.



NOTE: Prior to vGW Series Release 5.5, you used the Settings module Security Settings > Global Settings Rules pane to control broadcast and multicast settings. As of Release 5.5, you can no longer set these parameters from the Global Settings Rules pane. Rather, you must use the corresponding network object in a policy rule to control the firewall behavior.

Predefined Network Objects for Well Known IP Addresses

vGW Series provides the following predefined network objects that you can use in policy rule terms as either source or destination addresses:

- Link Local Addresses (**fe80::/10**)

IPv6 link-local addresses are defined in section 2.5.6 of the IETF RFC 4291 standard as having a 10-bit prefix of **fe80** followed by 54 zero bits and a 64-bit interface ID.

A link-local address is an IP address that is intended for communications within the link, or segment, of a local network or a point-to-point connection that a host is connected to. These addresses are useful for establishing communication across a link in the absence of a globally routable prefix or for intentionally limiting the scope of traffic that should not be routed. IPv6 link-local addresses, therefore, can be used only within the context of a single Layer 2 domain. Packets sourced from or destined to a link-local address are not forwarded out of the Layer 2 domain by routers.

- IPv4 Mapped Addresses (**::ffff:0.0.0.0 – ::ffff:255.255.255.255**)

The IETF RFC 6052 standard *IPv6 Addressing of IPv4/IPv6 Translators* covers the algorithmic translation of an IPv6 address to a corresponding IPv4 address, and vice versa, using statically configured information. Algorithmic translation is used in IPv4/IPv6 translators and other types of proxies and gateways that are used in IPv4/IPv6 scenarios, such as DNS.



NOTE: vGW Series accepts both IPv4 and IPv6 address formats and displays the addresses as you enter them.

- Well Known Prefix for IPv4 (**64:ff9b::/96**)

The IETF RFC 6052 standard *IPv6 Addressing of IPv4/IPv6 Translators* covers the Well Known Prefix **64:ff9b::/96** that is used in an algorithmic mapping between IPv4 to IPv6 addresses. It is defined out of the **0000::/8** address block.

- IPv4 Local Broadcast (**255.255.255.255**)

A special definition exists for the IP broadcast address **255.255.255.255**. It is the broadcast address of the zero network or **0.0.0.0**, which in IP standards implies the local network. Transmission to this address is never forwarded by the routers connecting the local network to other networks.

Additional IPv4 and IPv6 Predefined Network Objects

- Unspecified IPv4 (all zeros)

In IPv4, an IP address of all zeroes (**0.0.0.0**) has a special meaning. It refers to the host itself. It is used when a device does not know its own address.

- Unspecified IPv6 (all zeros)

The IPv6 unicast unspecified address is equivalent to the IPv4 unspecified address. The IPv6 unspecified address is **0:0:0:0:0:0:0:0**, or a double colon (::). In IPv6, this

concept has been formalized. It is typically used in the source field of a datagram sent by a device seeking to have its IP address configured.

- Loopback IPv4 (**127.0.0.1**)

The IETF RFC 2606 standard officially reserved domain name for the IPv4 and IPv6 loopback network addresses is localhost.

In IPv4, this network has the prefix **127.0/8**, as defined in the IETF RFC 3330 standard. The most commonly used IP address on the loopback device is **127.0.0.1** for IPv4, although any address in the range **127.0.0.0** to **127.255.255.255** is mapped to it.

- Loopback IPv6 (::1)

The IETF RFC 2606 standard officially reserved domain name for the IPv4 and IPv6 loopback addresses is localhost. IPv6 designates only a single address for the IP loopback function, ::1. The ::1/128 prefix is defined in the IETF RFC 3513 standard.

- Multicast IPv4 (**224.0.0.0/4**)

A multicast address is a logical identifier for a group of hosts in a network that are available to process datagrams or frames for a designated network service. IPv4 and IPv6 multicast addressing is used at Layer 3 (OSI) for IPv4 and IPv6.

The Classless Interdomain Routing (CIDR) prefix of multicast addresses is **224.0.0.0/4**. The group includes the addresses from **224.0.0.0** to **239.255.255.255**. Address assignments from within this range are specified in the RFC 5771 standard.

- Multicast IPv6 (**ff00::/8**)

Multicast addresses in IPv6 have the prefix **ff00::/8**. IPv6 multicast addresses are generally formed from 4-bit groups, illustrated as follows:

- Prefix: The **prefix** holds the binary value **11111111** for any multicast address.
- Flags: Currently, 3 of the 4 flag bits in the **flgs** field are defined. The left-most, most-significant flag bit is reserved for future use.
- Scope: IPv6 multicast addresses specify their scope. The set of possible scopes is different. The 4-bit **sc**, or scope, field (bits 12 to 15) is used to indicate whether the address is valid and unique.
- Group ID: The 112-bit **group ID** field identifies the service. For example, if **ff02::101** refers to all Network Time Protocol (NTP) servers on the local network segment, then **ff08::101** refers to all NTP servers in an organization's networks. The Group ID field may be further divided for special multicast address types.

Related Documentation

- [Understanding the vGW Series Firewall Module on page 93](#)
- [Configuring vGW Series Firewall Policies on page 114](#)
- [Understanding the vGW Series Policy per vNIC Feature on page 258](#)
- [Understanding vGW Series on page 3](#)

Configuring vGW Series Firewall Policies

This topic covers how to create a firewall policy for a VM composed of the corporate Global Policy, two Group Policies for the groups that the VM is a member of, and one VM Policy rule applicable to the individual VM.

It covers the preliminary tasks of defining the reusable Global Policy and a Group Policy for one of the groups that the VM is a member of.

Before you begin this procedure, read [“Understanding the vGW Series Firewall Module” on page 93](#). The procedure for composing an overall policy for a VM includes these parts:

- Define a Global Policy. The Global Policy is a reusable policy that is inherited by firewall policies for all VMs. You need to define it only once.

When you select the Firewall module and a VM in the VM tree to create a VM policy for it, the VM policy automatically inherits the Global Policy that you have created.

- Define Group Policies for the groups that the VM belongs to. You can define a Group Policy for a Policy Group any time after the Policy Group is created.

If the individual VM belongs to a Policy Group, it automatically inherits the Group Policy defined for that Policy Group, if the Group Policy is already defined.

When you select the Firewall module and a VM in the VM tree to create a VM Policy for it, the VM Policy contains the Group Policies that you created for any groups that the VM is a member of.

After you define the Group Policy for a group, it is automatically used in the individual policies that you construct for all members of the group. VMs that are created later and added to the policy group, either manually or automatically, inherit the Group Policy rules for that group.



NOTE: To illustrate precedence setting, this example assumes that the Group Policy already exists. It shows how to modify it.

- Define an individual VM Policy for the VM. At this point, you build the overall policy for the VM.

The VM Policy for a VM is composed of the Global Policy, Group Policies for any groups that it belongs to, and any individual VM Policy rules that you want to apply to that VM in particular.

When you select the Firewall module and a VM in the VM tree to create a VM Policy for it, the policy automatically inherits the Global Policy and the Group Policies for any groups that the VM is a member of. To complete the individual VM Policy, you add any VM Policy rules that you want to apply to that VM only. For example, you might need RADIUS access to a particular VM and not to others. You could apply a VM Policy rule to that VM's individual policy.

Create a reusable Global Policy to be used as part of the VM policies for all VMs in your environment.



NOTE: This example focuses on defining an inbound policy only. The process of defining outbound policy mirrors it.

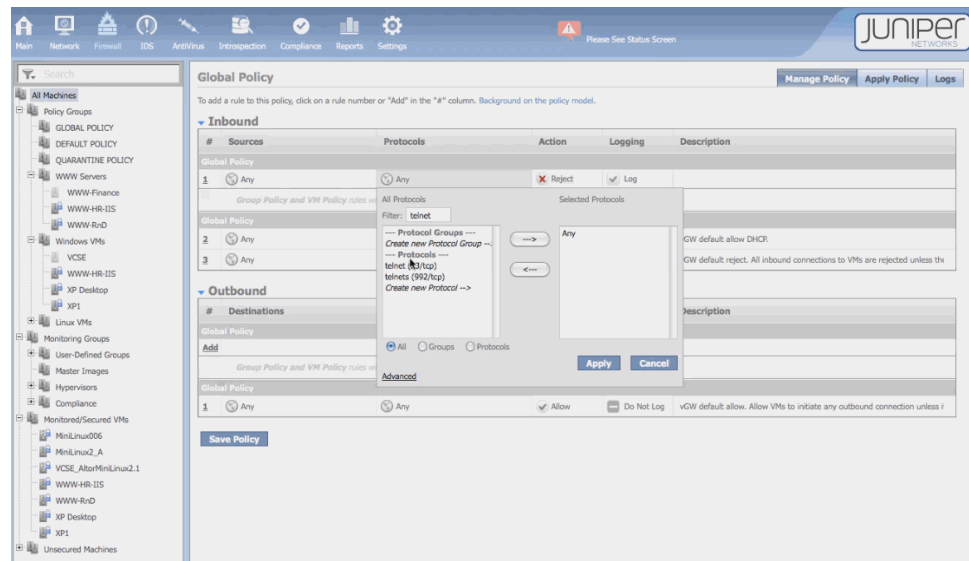
1. Define a Global Policy. From the Firewall module, select **Global Policy** under the Policy Groups section in the VM Tree.

The Global Policy page appears. It contains Inbound and Outbound sections. Each section contains a high-level Global Policy section and a low-level Global Policy section with a placeholder for Group Policy rules and individual VM Policy rules in the middle. [Figure 64 on page 115](#) shows the Global Policy with its default policy rules.

Figure 64: Default Global Policy

2. Create an Inbound high-level Global Policy rule to prohibit use of Telnet.
 - a. In the Inbound section, click **Add** in the # column under the first section labeled Global Policy to add a rule.
 - b. For the Sources policy term, leave the default value Any unchanged.
You want the rule to apply to all VMs.
 - c. Click **Any** in the Protocols column, and enter **telnet** in the Filter box. The filter scrolls to **telnet**.
 - d. Select **telnet**, and click the right arrow to move telnet from the All Protocols section to the Selected Protocols section. See [Figure 65 on page 116](#).

Figure 65: Adding a Global Policy Rule to Reject Telnet Connection Attempts



- e. Click **ALLOW** in the actions column and select **Reject** from the Action options list. You want to reject all inbound Telnet connections attempts for all VMs in your environment.
 - f. Leave the check mark default setting for Logging unchanged. Although they are rejected, you want to log any Telnet connection attempts.
3. Leave the low-level Global Policy rule unchanged.

By default, the last rule serves as a “clean-up” rule that catches all inbound connection attempts to this VM that have fallen through the rest of the policy rule base. It rejects them, and it specifies that vGW Series should create a log entry for the event.

4. Click **Save Policy**.

Modify the Group Policy for the Window VMs Policy Group to control rule execution precedence.

This procedure allows you to modify an existing Group Policy to change rule execution precedence. You want to ensure that a rule currently positioned in the low-level Group Policy section is not overridden by a VM Policy rule that might be inserted above it when an individual VM policy that includes the Group Policy is created. You want that rule to be executed *before* any VM Policy rules. To achieve that result, move the rule up from the low-level Group Policy section to the high-level Group Policy section.



NOTE: This example focuses on defining an Inbound policy only. An outbound policy definition process mirrors it.

1. In the Policy Groups section of the VM tree, select **Windows VMs**.

Notice that the high-level and low-level Group Policy sections are nested within the high-level and low-level Global Policy sections.

indicates the placeholder for adding VM Policy rules at the center of the Group Policy section.

2. Move the network management rule from the low-level Group Policy section to the high-level Group Policy section so that any VM Policy rule for an individual VM Policy rule added later cannot override it. See .
3. Click **Save Policy**.

Create a VM Policy for an individual VM

This procedure covers how to create individual VM policy rules for the WWW-HR-IIS VM that inherits the Global Policy and the Group Policies for the groups that it is a member of. An individual VM can belong to more than one Policy Group. When that is the case, the VM inherits the Group Policies for all of the Policy Groups that it belongs to. In this example, the WWW-HS-IIS VM is a member of two Policy Groups: WWW Servers and Windows VM.

This example focuses on the Inbound section of the VM Policy.

1. To display the VM Policy for the WWW-HR-IIS VM, select **WWW-HR-IIS** in the Windows VMs under Policy Groups in the VM Tree.



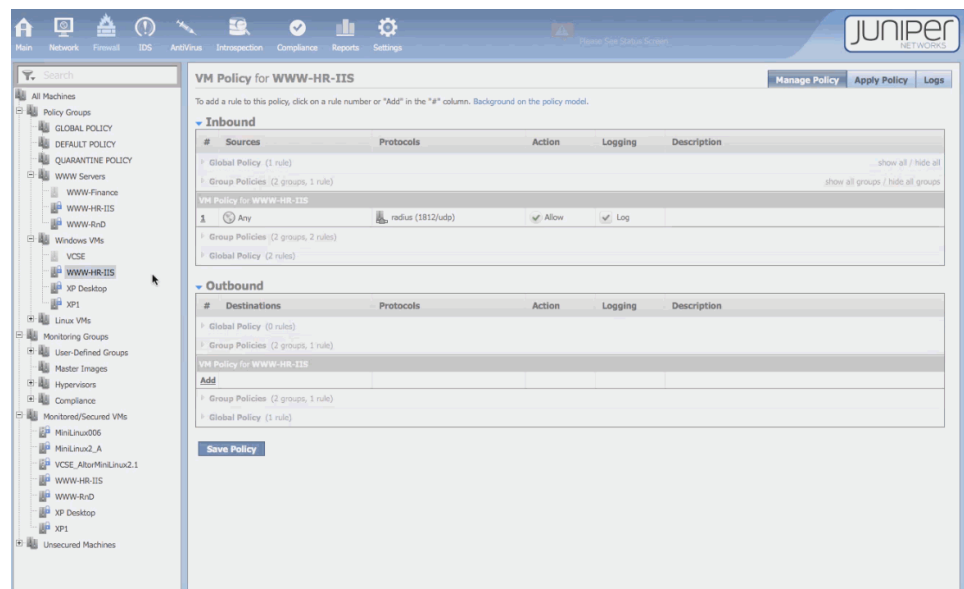
TIP: Because WWW-HR-IIS belongs to two groups, you can select it under either of its groups to display its VM Policy page.

The VM Policy for WWW-HR-IIS page is composed of the following nested parts that were previously built:

- the high-level and lower-level Global Policy rules forming the outer layer of the nest.
- a high-level Group Policy section below the high-level Global Policy. It states that the VM Policy contains two Policy Groups with a rule defined in only one of them.
- a middle section called VM Policy for WWW-HR-IIS. You can add VM Policies specifically for the VM to this section.
- the low-level Group Policy section that indicates that the VM belongs to two Policy Groups and that it inherits their Group Policies that include two rules.
- the low-level Global Policy.

Figure 66 on page 118 shows the policy.

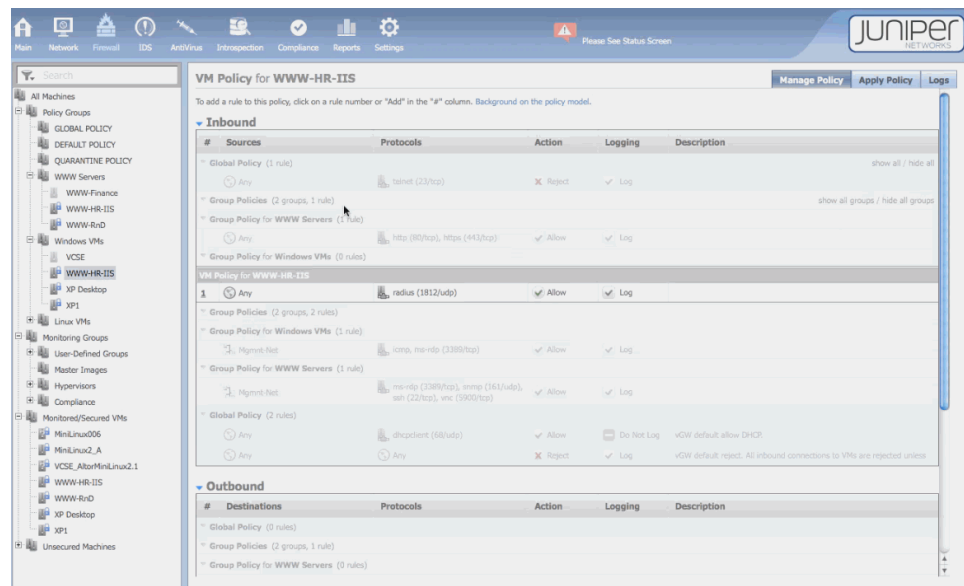
Figure 66: VM Policy for an Individual VM



- To see the entire rule base for the VM, expanding the policies that it inherited to show their rules, click **show all** in the upper-right corner of the page.

See [Figure 67 on page 118](#).


Figure 67: Complete VM Policy for an Individual VM




Apply the VM Policy.

When you define a firewall policy for a VM, it is not automatically applied. You must use the Firewall module Manage Policy tab to install it. This procedure installs a firewall policy for a single VM: AltorMiniLinux3.

1. Select the Firewall module. Select **All Machines** in the VM Tree. The following page is displayed.



2. Select the VM and click **Install**. In this example, All Machines is selected. After the firewall policy is installed on the VMs, the message shown in the following figure is displayed.



Related Documentation

- [Understanding vGW Series on page 3](#)
- [Using the vGW Series Network and Firewall Modules Cooperatively on page 90](#)
- [Understanding vGW Series Predefined Firewall Policy for Its Components on page 120](#)

Understanding vGW Series Predefined Firewall Policy for Its Components

vGW Series Firewall module allows you to secure virtual machines (VMs) within your virtualized infrastructure with individual policy rules, group policy rules, and global policy rules.

Not to be confused with securing VMs in your virtualized data centers, vGW Series secures and protects its own two main components—the vGW Security Design VM and the vGW Security VM—with predefined rule sets. You cannot change these predefined policy rules nor should you ever need to.

vGW Series stateful firewall comprises the following predefined rule sets for its two components.

For the vGW Security Design VM, the policy rules

- allow the following connections:
 - all outgoing connections
 - all incoming TCP/8443
 - all incoming TCP/443
 - all incoming TCP/8003
 - DHCP
 - NDP on IPv6
- Otherwise all connection attempts are dropped.

For the vGW Security VM, the policy rules

- allow the following connections:
 - all outgoing connections
 - all incoming TCP/8443
 - DHCP
 - NDP on IPv6
- Otherwise all connection attempts are dropped.

Related Documentation

- [Understanding the vGW Series Firewall Module on page 93](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding vGW Series on page 3](#)

CHAPTER 9

vGW Series IDS Module

This chapter includes the following topics:

- [Understanding the vGW Series IDS Module on page 121](#)
- [Configuring IDS Settings and Viewing Activity on page 128](#)

Understanding the vGW Series IDS Module

vGW Series includes a fully integrated IDS engine that you can use to monitor all virtual network traffic. It takes into account IPv4 and IPv6 traffic. You can also selectively monitor traffic for a subset of VMs or protocols used. vGW Series matches the selected traffic to the signature database and flags any suspicious activity with High, Medium, or Low priority alerts.

This topic covers the IDS module Alerts pages.

Use the Settings module > Security Settings > IDS Settings page to configure IDS for your environment. See [“Understanding and Configuring IDS Settings” on page 292](#).

The IDS engine shows attacks generated by VMs or by external systems. The IDS engine can identify an attack when one party involved in the attack is a VM.

This topic includes the following sections:

- [Managing and Sorting Displayed Alerts Information on page 121](#)
- [Top Alerts Page on page 122](#)
- [Alert Sources Page on page 127](#)
- [Alert Targets Page on page 127](#)
- [All Alerts Page on page 127](#)

Managing and Sorting Displayed Alerts Information

By default, basic alerts information is displayed for all Alerts tabs. In basic mode, you can change the time interval to control the period for which alerts information is displayed. Also, you can click the displayed information column heads to sort alerts based on alert type, signature ID, total number of alerts of that type, or priority.

For all Alerts tabs, advanced mode gives you the following additional capabilities:

- You can enable Auto-refresh to direct vGW Series to refresh, or update, the alerts information displayed every 60 seconds.
- You can direct vGW Series to sort displayed alerts information based on alert level. You might want to quickly view only high alerts to assess the greatest danger. In that case, you can select High and remove the check mark from the check boxes for Medium and Low alerts.

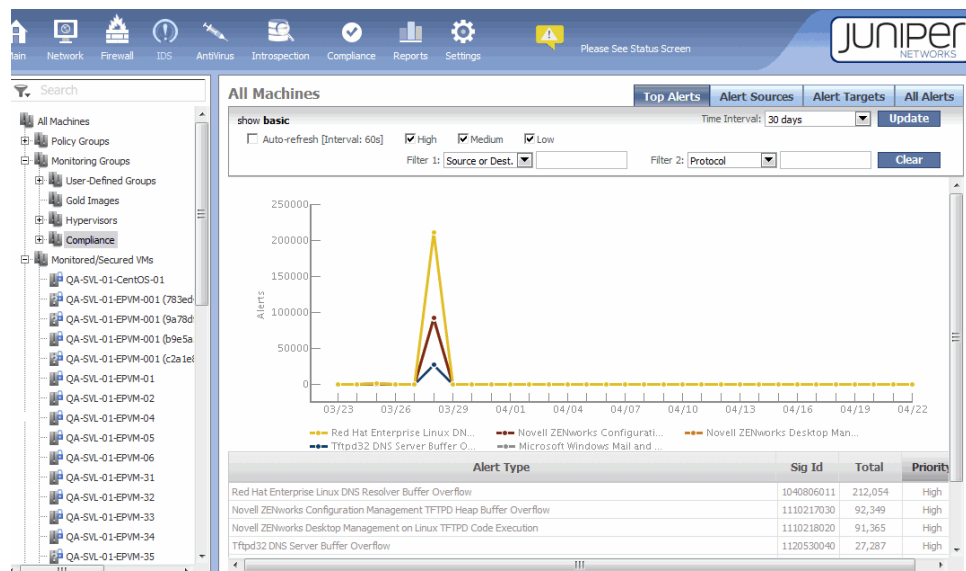
You can show all alerts—High, Medium, and Low alerts—as basic mode does, or you can show only High, Medium, or Low alerts, or any combination of them.

- You can use the advanced filter capability to display alerts information based on two filter settings. Your first filter can direct vGW Series to sort alerts based on Source or Destination, alerts for both of them, or by Signature ID, Protocol, or Record ID. Your second filter could refine even further the information that is displayed, specifying one of these categories in conjunction with the first filter value. For example, you might want to sort alerts by signature ID and within that result sort by Source to look at a specific kind of event and the sources that generated the alert.

Top Alerts Page

The Top Alerts tab presents a graph that shows the top alerts for attacks that have occurred over a specified period of time, for example 24 hours. If you specify a different time interval, alerts that have occurred within that period of time are displayed. The graph allows you to view at a glance for each alert type the degree of frequency. It includes a table that identifies the type of alert and its signature ID. See [Figure 68 on page 122](#)

Figure 68: IDS Top Alerts

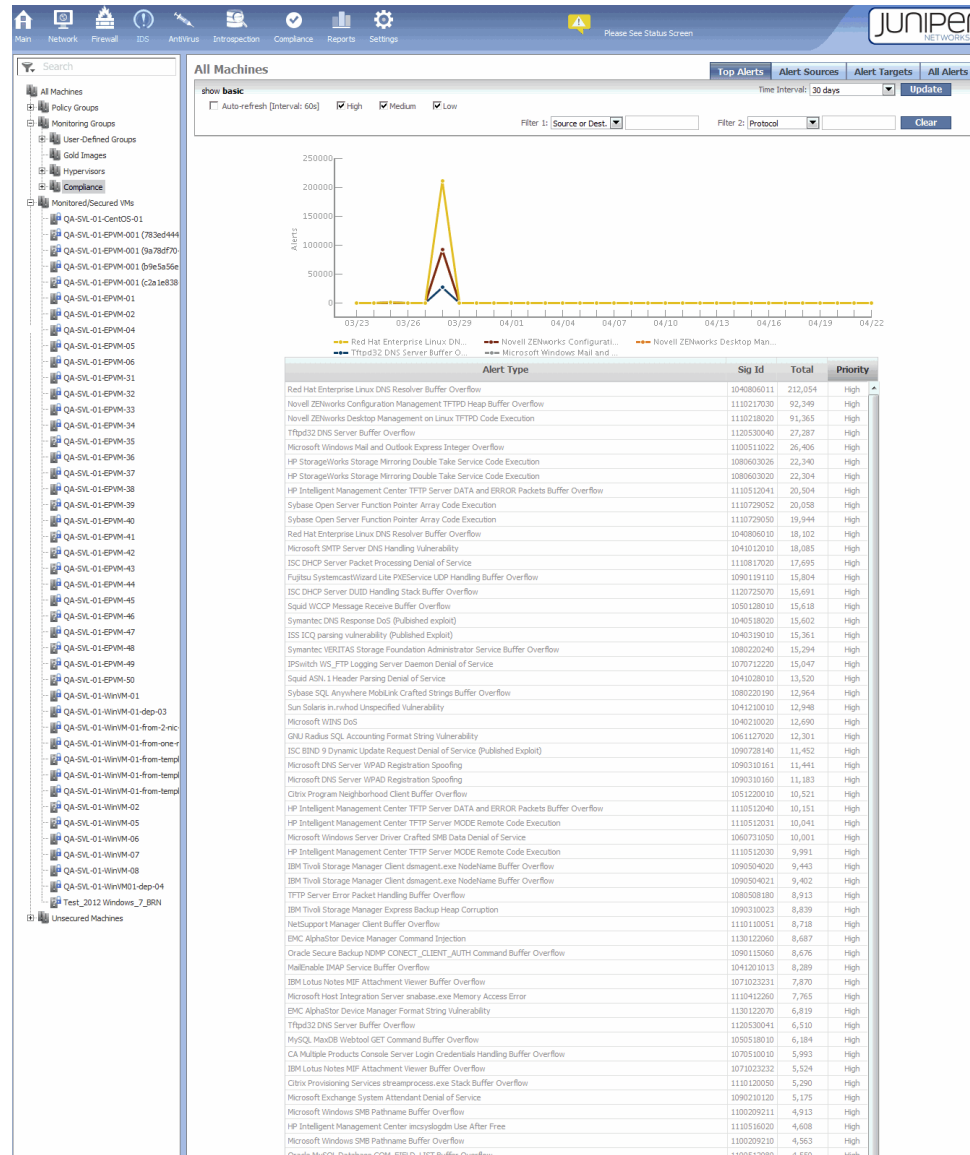


The alerts are organized as High, Medium, and Low with the total number sorting from most frequent to least frequent in the Total column.

To display advanced mode that gives you more options, click **show advanced**.

Figure 69 on page 123 shows the features that you can use in advanced mode with the time interval changed to reflect information for 30 days.

Figure 69: IDS Top Alerts Advanced Options



TIP: To change the priority level of an alert or not display information about it, use the Settings module > IDS Signatures page > Security Settings section.

To show information about a specific attack that caused the alert, click its row in the **Alert Type** column. In response, you see a description of the alert and its signature ID. See Figure 70 on page 124.

Figure 70: IDS Alert Description

The screenshot shows the Juniper Networks IDS interface. On the left is a navigation pane with categories like 'All Machines', 'Policy Groups', 'Monitoring Groups', 'User-Defined Groups', 'Gold Images', 'Hypervisors', 'Compliance', 'Monitored/Secured VMs', and 'Unsecured Machines'. The main area is titled 'All Machines' and displays a list of alerts. A modal window titled 'Alert Details' is open, showing information for the alert 'Tftpd32 DNS Server Buffer Overflow'. The alert has a priority of 'High'. The description reads: 'This rule checks to see if there is a non-zero number of DNS question records, and if so, checks the first domain name label length in the first question record, and if it is between 128 and 191 (inclusive), then an alert is generated.' Below the description is a 'show details' button. The 'Alert Details' box also shows the Sig Id: 1120530040 and references: Telus TSL20120530-04, OSVDB 82489, and Secunia SA49301.

To show additional details for that alert, beneath the alert description click **show details**. Figure 71 on page 124 shows the result.

Figure 71: IDS Alert Details

The screenshot shows the 'Alert Details' window for the 'Red Hat Enterprise Linux DNS Resolver Buffer Overflow' alert. The alert has a priority of 'High'. The description reads: 'This signature monitors DNS responses sent from port 53/UDP. If the Number of Answers value at offset 6 is greater than 48, an alert will be triggered.' Below the description is a 'hide details' button. The 'Alert Details' box also shows the Sig Id: 1040806011 and references: CVE CVE-2002-0029, Bugtraq 6186, and Telus TSL20040806-01.

Scroll down on the Alert Details box to see the affected systems and the attack scenarios. See Figure 72 on page 125.

Figure 72: IDS Alert Details Showing Affected Systems

Alert Details	
Alert Sources	Alert Targets All Alerts
Priority: High	
Red Hat Enterprise Linux DNS Resolver Buffer Overflow	
Description This signature monitors DNS responses sent from port 53/UDP. If the Number of Answers value at offset 6 is greater than 48, an alert will be triggered.	
hide details	
Affected Systems GNU C Library Project GNU C Library, version 2.3.1 and prior Red Hat Enterprise Linux, version AS 2.1 (glibc 2.2.4) Red Hat Enterprise Linux, version ES 2.1 (glibc 2.2.4) Red Hat Enterprise Linux Linux, version 6.0 (glibc 2.1.1) Red Hat Enterprise Linux Linux, version 6.1 (glibc 2.1.2) Red Hat Enterprise Linux Linux, version 6.2 (glibc-2.1.3) Red Hat Enterprise Linux Linux, version 7.0 (glibc 2.1.92) Red Hat Enterprise Linux Linux, version 7.1 (glibc 2.2.2) Red Hat Enterprise Linux Linux, version 7.2 (glibc 2.2.4) Red Hat Enterprise Linux Linux, version 7.3 (glibc-2.2.5) Red Hat Enterprise Linux Linux, version 8.0 (glibc-2.2.93) Red Hat Enterprise Linux Linux Linux Advanced Workstation, version Itanium 2.1 (glibc 2.2.4)	
Attack Scenarios	
Sig Id:	1040806011
References CVE-2002-0029 Bugtraq 6186 Telus TSL20040806-01	
Close	

If you want to know who generated the traffic that caused an alert, click the **Alert Sources** tab. See [Figure 73 on page 126](#).

Figure 73: IDS Alert Sources

IDS Updates

IDS signatures are updated frequently. The settings below control the behavior of the update processing.

Update Status

Currently Installed Signatures: **20130331021143**
 Signatures Available for Update: **20130428071231**
 Last Update Check: **Wed May 01 20:28:56 PDT 2013**
 Next Update Check:

[Check for Update](#) [Install](#)

Automatic Updates (Hourly Check)

☒ No Automatic Updates
☐ Download Automatically, Manually Apply Updates
☐ Download and Apply Update Automatically

[Save](#)

Manual Update

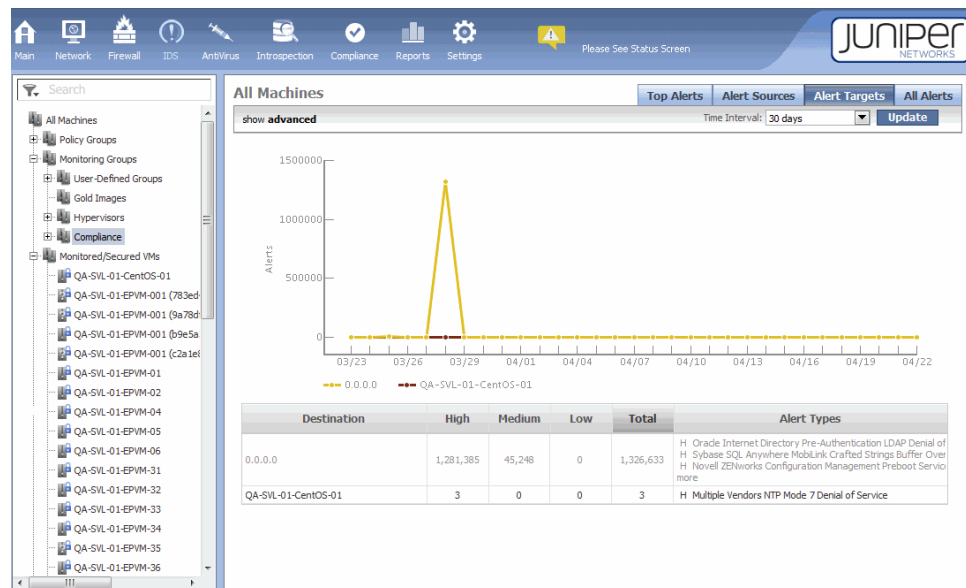
Manually upload an IDS signatures file for processing

[click to browse](#) [Browse...](#) [Clear](#)

[Upload File](#)

If you want to know the traffic destination, click the **Alert Targets** tab. See [Figure 74 on page 126](#).

Figure 74: IDS Alert Targets



Alert Sources Page

The Alert Sources window shows which systems have generated traffic matching the IDS signatures. These systems can be guest VMs or external systems communicating on the virtual network. The columns show High, Medium, and Low alert counts and a total count.

The system with the highest total count is displayed at the top of the list. You can sort the display by clicking the **High**, **Medium**, or **Low** columns. See [Figure 73 on page 126](#).

Alert Targets Page


The Alert Targets window shows the same information as the Alert Sources page but also it shows a list of the systems that are under the greatest number of attacks. See [Figure 74 on page 126](#).

All Alerts Page

The All Alerts tab shows a complete list of alerts for attacks captured by the system for the configured **time interval** (by default, 24 hours). In this example, the time interval has been set to 30 days.

To show details for a specific alert, click the alert type. By default, the most recent events are displayed at the top of the page, and older events are shown at the bottom. See [Figure 75 on page 127](#).

Figure 75: IDS All Alerts



The Source and Destination columns in the All Alerts page table show machine names, not IP addresses. When you roll the mouse and hover over a machine name, vGW Series displays its IP address. To make it clear which IP address is involved, vGW Series displays only the IP address that the alert pertains to, not all IP addresses for that machine.

Machines for which IPv6 is enabled typically have two addresses bound to each Virtual Network Interface Card (vNIC)—a link local address and a routable address. Typically the link-local address is not used by applications. A machine can have multiple vNICs, each of which might have two IP addresses. Effectively a machine might have many IP addresses bound to it.

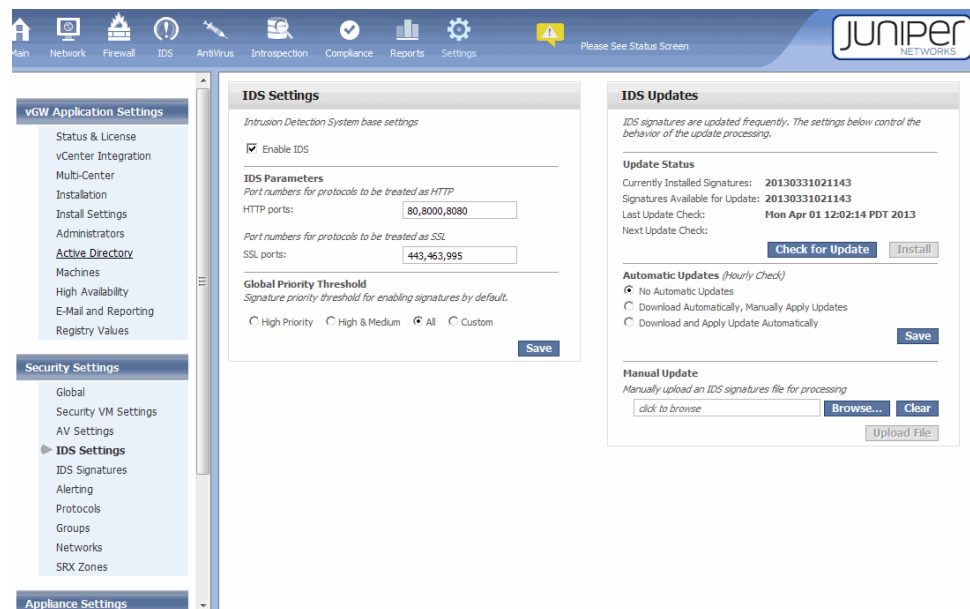
- Related Documentation**
- [Understanding and Configuring IDS Signatures Settings on page 295](#)
 - [Configuring IDS Settings and Viewing Activity on page 128](#)

Configuring IDS Settings and Viewing Activity

This topic covers how to configure IDS and view the results produced by the IDS engine.

1. Enable IDS and specify its settings using the Settings module Security Settings > IDS Settings > IDS Settings pane. See [Figure 76 on page 128](#).

Figure 76: IDS Settings Page



2. Enable the signatures relative to your environment.

From the Settings module, select Security Settings > IDS Signatures for a list of signatures.

For details, see [“Understanding and Configuring IDS Signatures Settings” on page 295](#)

3. Create and apply a policy rule that mirrors traffic to the IDS engine. vGW Series gives you the ability to specify at a granular level which traffic to scan. For example, you might want to scan traffic to or from a specific VM, or traffic that uses a specific protocol.



NOTE: Traffic that the firewall blocks is not inspected by the IDS engine because the connection is never established.

A policy rule might be defined to inspect a connection for IDS but that does not imply that it accepts it. If the policy rule accepts, drops, or rejects a connection—all of which are considered terminal actions—policy scanning terminates. In this case, IDS rules that follow the rule that caused policy scanning to terminate are not processed. For IDS to take effect, the IDS rule for a connection must precede the rule that accepts the connection.

**Related
Documentation**

- [Understanding the vGW Series IDS Module on page 121](#)
- [About the vGW Series IDS Reports on page 200](#)
- [Understanding and Configuring IDS Signatures Settings on page 295](#)

CHAPTER 10

vGW Series AntiVirus Module

This chapter includes the following topics:

- [Understanding vGW Series AntiVirus on page 131](#)
- [vGW AntiVirus Configuration Overview on page 139](#)
- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Understanding Quarantined VMs and Files Resulting from a vGW AntiVirus On-Access Scan on page 148](#)
- [Understanding and Installing the vGW Endpoint on page 149](#)
- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)
- [Understanding Quarantined VMs and How to Manage Them on page 157](#)

Understanding vGW Series AntiVirus

This topic explains the vGW AntiVirus feature. vGW AntiVirus provides improved security and flexibility that agents alone cannot provide. It does this through:

- use of its kernel module installed in the ESX/ESXi host hypervisor.
- its management integration.
- its On-Access scans on VMs with only a light installation on the machine using its vGW Endpoint. An on-access scan is performed whenever a file is read from or written to disk.
- its On-Demand scans on VMs entirely without any installation on the VM and including no requirement to reconfigure the VM after the scan. vGW Series takes snapshot of the VM disk, and it performs the scan offline and deletes the snapshot that it takes and scans.

This topic begins by giving background information on antivirus technology. Then it explains vGW AntiVirus.



NOTE: vGW Series AntiVirus feature requires a license.

For an overview of the complete vGW AntiVirus configuration process, including information on mandatory preliminary configurations, read "[vGW AntiVirus Configuration](#)"

[Overview](#)” on page 139. For each step, the topic provides links to topics that give detailed procedures.

This topic includes the following sections:

- [About Antivirus Software](#) on page 132
- [Signature-Based Detection](#) on page 132
- [The vGW AntiVirus Feature](#) on page 133

About Antivirus Software

Antivirus software prevents and detects malware, such as viruses, worms, and spyware. A variety of strategies are usually involved in implementing antivirus software, including use of signature-based detection and rootkit detection, both of which the vGW AntiVirus supports.

Virtualized environments experience the same persistent threats and proliferation of malware that physical networks do. Not uncommonly, administrators of physical networks who have virtualized their environments install the same antivirus software that they use on their hardware desktops on their virtual machines. When it is installed on virtual systems, antivirus software designed for physical environments is severely limited, and it creates many problems. It does not recognize the virtual infrastructure; it consumes excessive memory usage, often exceeding 100 MB of RAM for a single guest VM; and it heavily degrades system performance through exhaustive CPU usage, often resulting in what is referred to as *brownout*.

Antivirus software is often the first line of defense against malware, but it should not provide this protection in the virtual environment at the cost of system performance.

Signature-Based Detection

A signature is a unique string of bits, or a byte pattern, that is characteristic and part of a certain virus or group of viruses. During a virus scan, the vGW AntiVirus feature compares the content of resources and files to be scanned against its virus signature database.

When vGW Series detects a signature pattern, it takes the remediation action that you specify when you configure the vGW AntiVirus scan. You use the vGW AntiVirus module's Scanner Config tab, which allows you to specify more than one action, for this configuration.

For example, when you select **Alert when a virus is detected** as an action, the Virus Alerts tab shows details on the event when vGW AntiVirus detects a virus. You can view the Virus Alerts tab content to gain an understanding of the types of threats that have been found, such as worm.exe, and where the threat was identified, such as the workstation name and other related information.

The vGW AntiVirus feature is robust in that it uses two methods to detect viruses and malware. It uses a signature database to detect specific viruses. It complements this approach with heuristics methods for detecting suspicious code parts.

The vGW AntiVirus Feature

Traditionally and extending into the present, antivirus software for the physical environment was developed to protect either the host—your desktop, servers, and other local devices—or the network for which malware and attack attempts could be caught before they reached the host.

Software for the desktop, and other hosts, is thought of as agent, or endpoint, software. Endpoint software involved installing a scanning engine and an attack signature database on every machine, which results in slower system startup and performance on the device. When device scans run, memory is consumed and performance is affected. This model was carried into the virtualized environment as security products began to become available for it; the virtualized network and the virtualized host were protected separately by separate products.

The vGW AntiVirus feature constrains performance impact on the VM in both cases by centralizing its scanning engine and signature database on the vGW Security VM firewall instantiated on each ESX/ESXi host for which you configure vGW AntiVirus, and not on each VM. For On-Access scanning, whenever a VM's disk is written to or read from, the “lightweight” vGW Endpoint that you install on it passes several portions of the file necessary to determine if it contains a virus to the vGW Security VM across the virtualized network for examination.

The vGW AntiVirus feature remains effective when VMware VMotion is used. When a VM that is protected by vGW AntiVirus is migrated to another ESX/ESXi host through VMotion, the VM remains protected. The vGW Security VM on the host to which it is moved takes up the vGW AntiVirus protection work, based on the original configuration.

The vGW AntiVirus feature protects VMs by detecting malware, quarantining affected VMs and for On-Access scans also quarantining affected files. It allows you to define a remediation plan.

When you enable the vGW AntiVirus feature, the vGW Security Design VM activates its scanning engine on the vGW Security VM. This approach centralizes the scanning engine to limit disk, disk I/O, memory, and CPU consumption, and distribute the load across the virtualized infrastructure. The vGW AntiVirus database and the updates to it are also deployed on the vGW Security VM.

vGW AntiVirus relies on three main components:

- vGW Security Design VM

You use the vGW Security Design VM to enable vGW AntiVirus, configure scans, view reports and alerts, download new signature versions, and download the vGW Endpoint.

If the vGW Security Design VM is configured for dual stack, first it attempts to use the IPv4 protocol to communicate with the vGW Security VM.



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

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

By default, this parameter is set to true.

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

- vGW Security VM

The vGW Security VM performs On-Demand scans.

It is possible to perform an On-Demand scan on a VM whose ESX/ESXi host does not have a vGW Security VM installed. In this case, the scan is performed by the vGW Security Design VM, a vGW Security VM on a different host (TCP 902 is required), or both.

vGW AntiVirus remains in effect when a VM is VMotioned to another host for analysis. In that case, the vGW Security VM on that host performs the vGW AntiVirus functions.

- vGW Endpoint

The vGW Endpoint is used for On-Access scans. It protects a VM against infected files whenever a file is read from or written to disk. The vGW Endpoint sends the file to the vGW Security VM to be analyzed.

When an infected file is identified and the quarantine action is specified in the On-Access scanner configuration, the file is isolated in the vGW Endpoint on the VM. It remains there until you *un-quarantine* it, delete it, or fetch it. When you release it from quarantine, it is made available to the VM again.



NOTE: On-Demand scans do not require installation of the vGW Endpoint. The vGW Endpoint is used for On-Access scans only.

vGW Series supports both AntiVirus On-Demand and On-Access features in IPv4 or IPv6 environments, or environments that are a mix of the two.

Although the vGW AntiVirus works in an IPv6 environment, communication between the vGW Endpoint and the vGW kernel module installed in the ESX/ESXi host hypervisor occurs over the IPv4 infrastructure. Note that the vGW Endpoint OS should be configured with the IPv4 stack enabled.

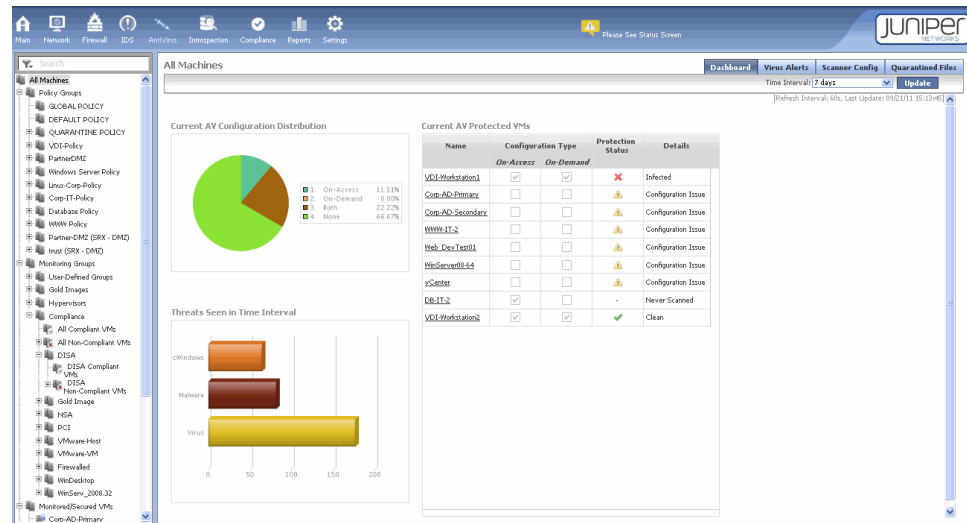
The vGW AntiVirus Dashboard

The vGW AntiVirus dashboard gives you an overall view of the current state of all protected VMs in your environment.

- You can view information for all VMs in your environment or for specific VMs. You use the VM tree to select the VMs.
- You can change the time interval to view threats that occurred within a broad or narrow span of time.
- You can view information on vGW AntiVirus events for VMs, such as details on viruses that were detected and signature updates.

Figure 77 on page 135 shows the vGW AntiVirus Dashboard.

Figure 77: vGW AntiVirus Dashboard




The vGW AntiVirus Dashboard includes these panes:

- Current vGW AntiVirus Configuration Distribution

This pie chart shows you proportionally the number of VMs that are protected by the On-Access scanner, by the On-Demand scanner, or both of them, and those that are not protected by vGW AntiVirus.


Current AV Configuration Distribution



- Threats Seen in Time Interval

This bar graph displays the kinds and percentage of threats that were identified in the selected time interval.

Threats Seen in Time Interval

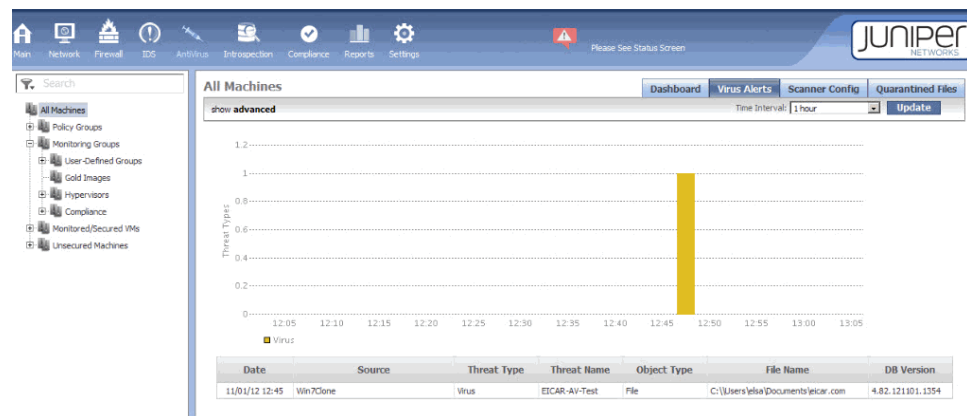


- Current vGW AntiVirus Protected VMs

This table identifies VMs that are protected by vGW AntiVirus, the type of scanner configurations that protect them, and the protection status and details for the VM. If the protection status indicates problems, you can click the VM's row to display a page dedicated to it giving detailed information. The page shows scan statistics for the VM (how many files were scanned, how many files were quarantined, and so on), the scanner configuration for the VM, the threat type bar graph as applied to the VM, and a table identifying attempted virus infections, when they occurred, and how vGW AntiVirus handled them.

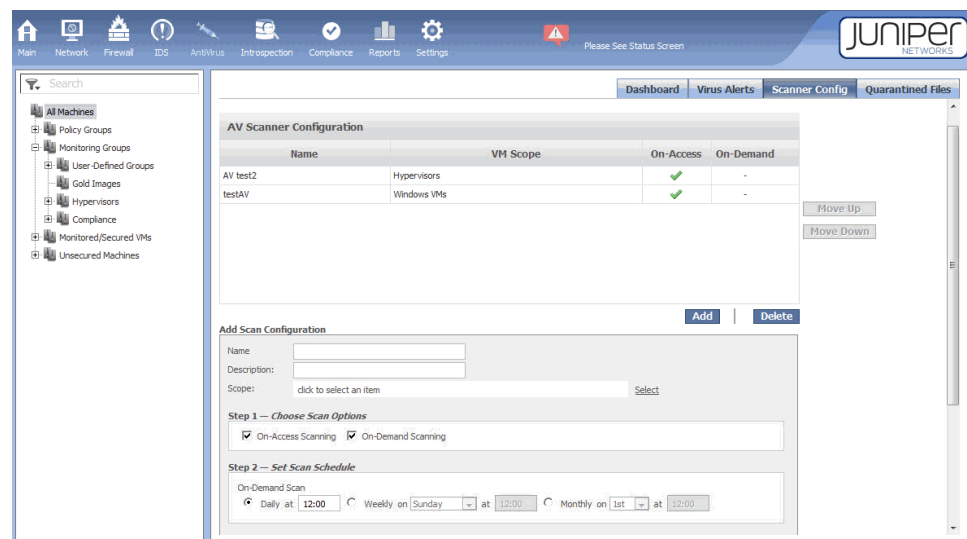
The Virus Alerts tab displays a graph that identifies threat types over a period of time. You use the Time Interval box to control the period. It gives details on the threat type, including the date of the event, the source, and the filename.

Figure 78: Virus Alerts



The Scanner Config page allows you to define On-Access and On-Demand scans. When you click **Add** to display the Add Scan Configuration pane, both types of scans are selected. You can configure them separately or together in one configuration. You can configure a typical scan or a custom scan. [Figure 79 on page 137](#) shows them configured together by default with a typical scan used. For details on configuring them separately, see “Configuring vGW Series AntiVirus On-Access Scanning” on page 145 and “Configuring vGW Series AntiVirus On-Demand Scanning” on page 153.

Figure 79: vGW AntiVirus Scanner Config Tab

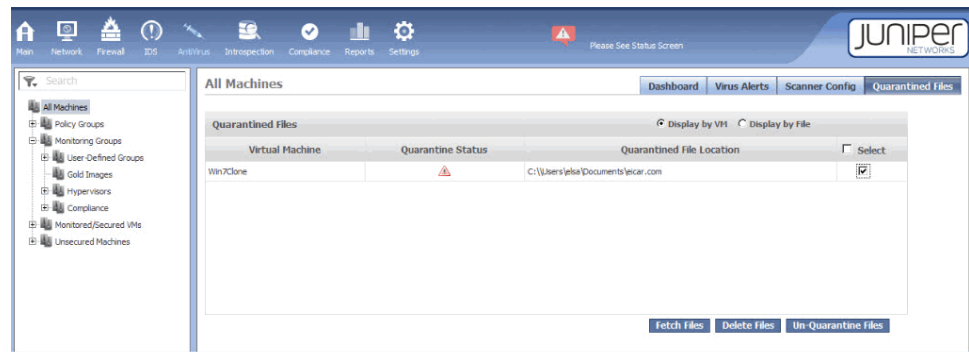


The Quarantined Files tab displays a list of quarantined files. Only infected files identified through an On-Access scan can be quarantined. When a file is quarantined, it is isolated in the vGW Endpoint on the VM and information about it is displayed on this page. The VM containing the file is identified. The location of the file is shown and its status is noted. See [Figure 80 on page 138](#).



NOTE: There must be no items for a VM in quarantine for that VM to appear as non-infected, or in a “clean” state, on the dashboard. However, if a VM is not quarantined and none of its items are quarantined does not mean that the VM is clean. If a VM has items in quarantine is not considered clean.

Figure 80: Quarantined Files



You can select one or more files and perform any of the following actions:

- You can fetch the file. In this case, the file is hashed and transferred off the VM for further analysis.
- You can un-quarantine the file. In this case, the isolated file is made available again to the VM.

In some cases, files are quarantined because of false positive results. That is, the file is suspected of being malware or infected, but that is not the case. Updating the signature database and running the scan again often resolves the problem.

- You can delete the file from the VM if you have confirmed that the file is infected or that it is malware.

When a VM is infected by a virus and the scanning configuration specifies **Quarantine the VM**, the VM is put in the quarantine policy group. To remove the VM from the quarantine policy group, use the Main module Quarantine tab. Select the VM, and click **Un-quarantine**.

For details on how the parts of the quarantine process work together for a quarantined VM, see [“Understanding Quarantined VMs and How to Manage Them” on page 157](#).

Related Documentation

- [Understanding and Installing the vGW Endpoint on page 149](#)
- [Understanding vGW Series on page 3](#)
- [Understanding and Configuring the vGW Series AntiVirus Settings on page 291](#)
- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)

vGW AntiVirus Configuration Overview

This topic gives an overview of the steps to follow to configure vGW AntiVirus protection for your virtualized environment.



NOTE: The vGW AntiVirus feature requires a license.

For vGW Series to scan a VM, the VM must be included in one of the VM groups that you include in the scan scope, which you define when you configure a scan. You use the AntiVirus module Scanner Config page to configure scans. If a VM is not included in one of the groups in the scope, it will not be protected by vGW AntiVirus. You can define at a granular level the files on a VM to be scanned based on file type and file location. For example, you can configure a scan to scan all file types, only certain file types, files at all locations or only files at certain locations. You can combine these options, for example, to scan all file types but only at a certain location. You can also refine the scan by excluding types of files or files at certain locations from it.

vGW AntiVirus provides two means of protecting your environment against malware and viruses:

- The On-Access Scanner


To protect VMs against malicious content and virus infections in real-time, the On-Access scanner runs whenever a VM's disk is written to or read from. It scans areas of the disk based on your configuration.

- When you configure a custom On-Access scan, you can specify file types and the location of files to scan. You can also exclude certain types of files and files at certain locations from the scan.
- You must specify the drive when you specify the location of files to scan for custom On-Access scans. For On-Access scans, when scanning files based on location, vGW Series takes into account the drive letter of the directory. For example, given the file location C:\Program, an On-Access scan scans files only in that directory. It does not scan files in the D:\Program directory, although the directory names are the same, because it acknowledges that the drive letters are different.
- For On-Access scans, vGW Series does not support the use of wildcards in file extensions or file locations.

Here is how the vGW AntiVirus On-Access scanner works, as illustrated in [Figure 81 on page 140](#).

Figure 81: On-Access Scan

ESX or ESXi host



1. vGW Series installs a small agent called the vGW Endpoint on the VM when On-Access scanning is configured.
2. The vGW Endpoint captures file accesses and forwards them to the vGW Security VM on the host to scan.

The file transfer is controlled internally, based on its match against the AntiVirus signatures. Only as much of the file as is necessary to determine if it is malicious is forwarded to the vGW Security VM. The vGW Security VM scans the file to make the determination. You cannot control this from the vGW Security Design VM.

3. An On-Access AntiVirus scan is performed.

The vGW Security VM scans the file to make the determination.

Because the scan is performed on the vGW Security VM, it is not necessary to re-configure a VM after an On-Access scan.

4. The scan results are cached in the vGW Endpoint for improved performance.

For details on how to configure On-Access scanner, see ["Configuring vGW Series AntiVirus On-Access Scanning" on page 145](#).

- The On-Demand Scanner

The On-Demand scanner performs full-disk offline scanning that scans VMs periodically, examining their virtual disk files for malicious content. You configure a schedule to specify when scanning should occur.

On-Demand scans are performed without any impact to the VM. The scanning is done outside the VM on the ESX/ESXi host's vGW Security VM. Therefore it not necessary to re-configure a VM after an On-Demand scan.

Here is how the vGW Series On-Demand scanner works:

1. vGW Series takes a snapshot of the VM disk to be scanned.
2. It attaches the snapshot to the vGW Security VM.
3. Based on your Scanner Config for On-Demand scans, it performs either a typical scan or a custom scan. For a custom scan, it scans the archives, file types and file locations that you specify, excluding any file types or locations that you specify in your custom scan configuration.
4. After it completes the scan, vGW Series detaches the snapshot from the vGW Security VM.
5. Finally, it deletes the snapshot.

Take into account the following characteristics when you configure a custom On-Demand scan:

- vGW Series recognizes the global wildcards * and ?.
For example, you could specify C:\Program Files\MS*. You could also use the wildcard on an extension, for example doc*.
- For file locations, drive letters are ignored. For example, C:\Program Files matches: C:\Program Files and D:\Program
vGW Series performs an On-Demand scan offline and does not take into account drive letters.

When you configure the vGW AntiVirus scanner, you can specify the action to take in response to results of the scan. Both On-Access and On-Demand scanning can result in a quarantined VM. However, files can be quarantined only as a result of an On-Access scan.

You can configure both On-Access Scanning and On-Demand Scanning in a single vGW AntiVirus scanner configuration.

You use the vGW AntiVirus module tabs in concert:

- to gain an overall, quick status on your environment as it stands in relation to vGW AntiVirus protection.
- to enact scanning.

- to identify files for which there are issues that need to be addressed and files that are quarantined.

For details on how quarantined VMs are treated, see [“Understanding Quarantined VMs and How to Manage Them”](#) on page 157.

Figure 82 on page 142 shows the vGW AntiVirus dashboard that gives you a comprehensive view of vGW AntiVirus protection for your environment. It emphasizes a table that shows vGW AntiVirus details on individual VMs, including the kind of vGW AntiVirus protection it has and the current scan status on the VM. The dashboard also presents a pie chart that shows the vGW AntiVirus protection distribution across VMs. It includes a chart that shows the types and degrees of threats identified by vGW AntiVirus across a specific period of time, which you can adjust.



NOTE: There must be no items for a VM in quarantine for that VM to appear as non-infected, that is, in a “clean” state, on the dashboard. However, simply because a VM is not quarantined and none of its items are quarantined does not mean that the VM is clean. But you can be assured that it is never the case that a VM that has items in quarantine is clean.

Figure 82: vGW AntiVirus Dashboard

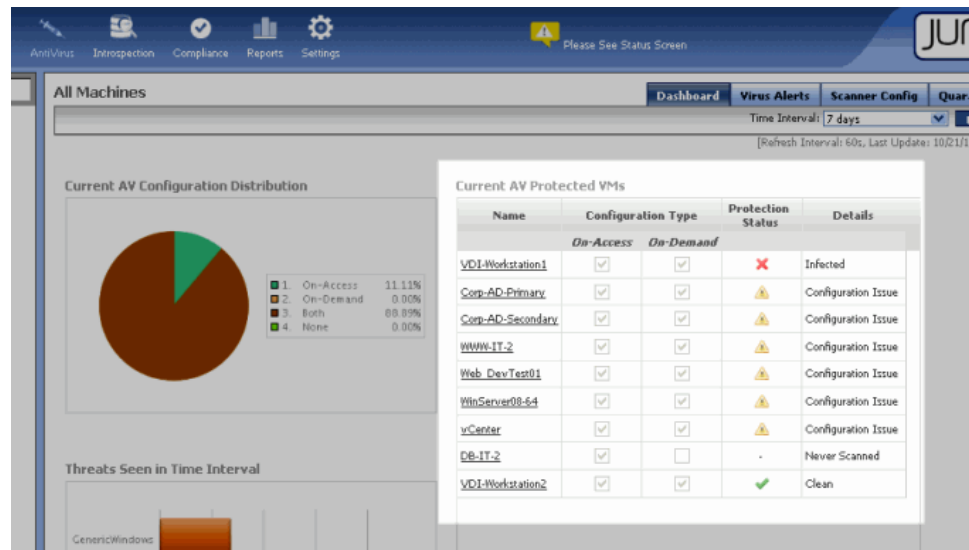
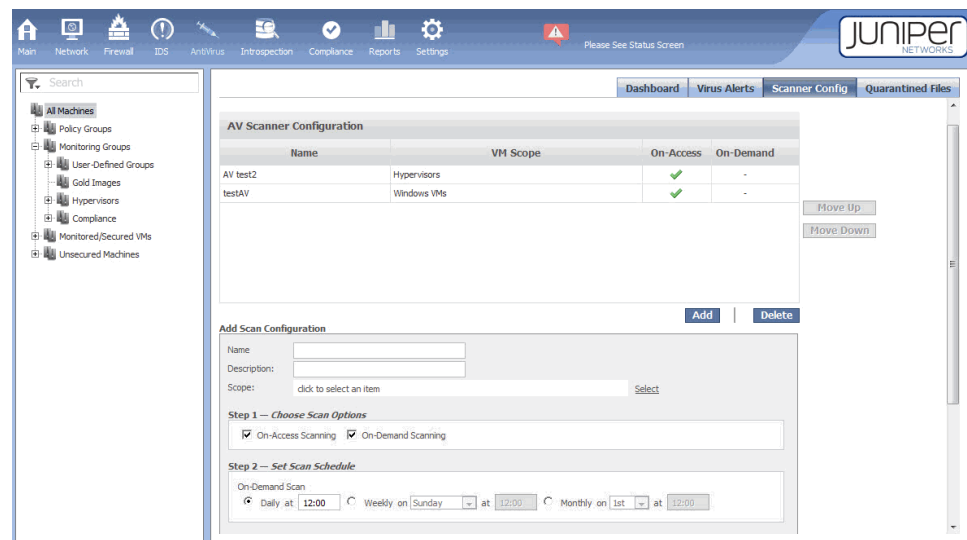


Figure 83 on page 143 shows the two scanning options that you can configure using the Scanner Config tab.

Figure 83: Scanner Config Tab



A vGW AntiVirus On-Access scan can result in quarantined files or VMs:

- Quarantined files are identified in the vGW AntiVirus module Quarantine tab.
- Quarantined VMs are identified in the vGW Main module Quarantine tab.

Complete these prerequisite tasks:

1. Secure the ESX/ESXi hosts. Deploy the vGW Security VM out to the ESX/ESXi hosts in your environment. From the Settings module, select **vGW Application Settings > Installation** for this purpose. See [“Installing vGW Security VMs on ESX/ESXi Hosts” on page 218](#).

If you do not deploy the vGW Security VM and you protect the VMs with the vGW firewall, On-Access scanning will not work. Configuring only the On-Access scanner for the VMs and enabling vGW AntiVirus is ineffective without this preliminary configuration.

2. Secure the VMs. Configure the vGW Firewall for VMs that you want to protect with On-Access scanning. From the Firewall module, select the **Manage Policy** tab to create firewall policies and the **Apply Policy** tab to apply them. See [“Understanding the vGW Series Firewall Module” on page 93](#).

To configure vGW Series On-Access scanning for your environment, you must:

1. Create an On-Access scanner configuration for the VMs.

See [“Configuring vGW Series AntiVirus On-Access Scanning” on page 145](#).



NOTE: When you configure an On-Access scan, you do not configure a scanner schedule. On-Access scanning occurs in real time.

2. Enable the vGW AntiVirus feature and download the vGW Endpoint.

See [“Understanding and Configuring the vGW Series AntiVirus Settings” on page 291](#).

3. Install the vGW Endpoint on the VMs to be protected.

See [“Understanding and Installing the vGW Endpoint” on page 149](#). This topic explains how to install the vGW Endpoint on VMs, and it explains the pop-ups that the vGW Endpoint displays to inform you about various conditions, such as when a threat is detected.



NOTE: You must install the vGW Endpoint on all VMs that you want to protect with On-Access scanning.

On-Demand scanning differs from On-Access scanning in the following ways:

- It is not possible to quarantine files when On-Demand scanning is used.
- You can run an On-Demand scan on VMs whose ESX/ESXi is not protected by the vGW Security VM. In this case, the scan is performed by the vGW Security Design VM, a vGW Security VM on a different host, in which case TCP 902 is required, or both.
- You do not need to install the vGW Endpoint on the VMs.
- For On-Demand scanning, you can protect VMs to be scanned with the vGW Firewall, but it is not required.

Because you do not need to protect VMs with the vGW Firewall and you do not need to install the vGW Endpoint on the VM, On-Demand scans can be performed on virtual disk files from a protected location that is not compromised. This advantage increases the ability of the vGW Series to detect and locate rootkits. It can detect files with suspicious names such as mal.exe, simpletroj.exe, and other malware files.

To configure On-Demand scanning:

1. Create an On-Demand scanner configuration for the VMs.

See [“Configuring vGW Series AntiVirus On-Demand Scanning” on page 153](#).

2. Enable the vGW AntiVirus feature.

See [“Understanding and Configuring the vGW Series AntiVirus Settings” on page 291](#).

Related Documentation

- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Understanding and Configuring the vGW Series AntiVirus Settings on page 291](#)
- [Understanding and Installing the vGW Endpoint on page 149](#)
- [Understanding vGW Series AntiVirus on page 131](#)
- [Understanding Quarantined VMs and Files Resulting from a vGW AntiVirus On-Access Scan on page 148](#)
- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)
- [Understanding the vGW Security VM on page 68](#)

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

Configuring vGW Series AntiVirus On-Access Scanning

This topic explains how to configure a vGW AntiVirus On-Access scanner configuration using the AntiVirus module Scanner Config tab. The On-Access scan protects VMs against malicious content and virus infections that can occur whenever a file is read from or written to disk. If On-Access scanning is configured, vGW AntiVirus intercedes and checks the file against the signature database to ensure that the content does not contain malware or a virus. By blocking an infected file, On-Access scanning protects the network from malicious attacks at the source, before damage is done.

Before you configure a vGW AntiVirus On-Access scan, you must perform prerequisite tasks. These tasks configure other parts of the system that allow vGW AntiVirus to quarantine an entire VM with the Quarantine policy when the VM is compromised by a virus. They also initiate communication with the vGW Endpoint:

- Deploy the vGW Security VM to the ESX/ESXi hosts in your environment.
- Configure and install firewall policies on the VMs to be protected.

When you configure a custom On-Access scan, you can specify types of files and files at certain locations to be scanned, and you can exclude certain types of files and files at certain locations from the scan. You can combine these options, for example, to scan all file types but only in a certain directory or to exclude certain types of files in a certain directory.

Consider the following characteristics, when you configure custom On-Access scans:

- When specifying a file location, for On-Access scans you must always specify the drive letter.
- vGW Series does not support the use of wildcards in specifying file types or file locations.

The vGW Endpoint captures file accesses and forwards them to the vGW Security VM for analysis. The vGW Endpoint driver caches the results of the scan. You cannot control how much of a file is transferred to the vGW Security Design VM. However, the file transfer, which is controlled internally, is efficient, based on its match against the AntiVirus signatures. Only as much of the file as is necessary to determine if it is malicious is forwarded.



NOTE: Because the scan is performed on the vGW Security VM, it is not necessary to re-configure a VM after an On-Access scan.

To create an On-Access vGW AntiVirus configuration or add a new one:

1. Select the AntiVirus module **Scanner Config** page.

The AV Scanner Configuration table is displayed showing information about existing AV scanner configurations. The table shows the scanner configuration name, the scope of VMs that the scan covers, and the type of scan: On-Access, On-Demand, or both.

2. Click **Add**.
3. Specify a name for the AntiVirus scanner configuration.
4. (Optional). Give a brief description of the scanner configuration so that it is quickly recognizable.
5. In the Scope box, identify the VM groups whose VM members are to be scanned.

For a VM to be protected by vGW AntiVirus, it must belong to a VM group that you include in the scan scope.

To select the scope, click **Select**. A pop-up dialog box is displayed that shows all VMs groups on the left side. Click on the name of a group and move it to the **Selected Groups** section on the right. Click **Apply**.

After a scan is defined, it is added to the list of configurations in the AV Scanner Configuration table.



NOTE: If a VM group is a member of more than one scanner configuration, the topmost scan definition that it belongs to is used to protect it. You can manipulate the order of the scanner configurations in the table by selecting the row for the scanner configuration and clicking either **Move Up** or **Move Down**.

6. In the **Step 1 Scan Options** pane, select the **On-Access Scanning** check box. By default, both types of scans are selected. In this case, clear the check box for **On-Demand Scanning**.



NOTE: Step 2 in the scanner configuration page is required for On-Demand scans only, so it is not included in this procedure.

7. In the **Step 3 Configure Scanning Engine** pane, select the type of scan to perform. Under **On-Access file types/extensions scanning selection**, select either **Typical Scan** or **Custom Scan**. For this example, select the **Typical Scan** check box.
8. In the **Step 4 Action** pane, specify one or more actions to take when the scan detects a virus:
 - **Alert when a virus is detected**—The Virus Alerts tab displays information on the VMs or files that are infected.
 - **Quarantine VM**—You can specify that the infected VM is to be included in a quarantine policy group.

You use the Quarantine page on the Main module to view a list of VMs quarantined as a result of an AntiVirus scan. From the Main module Quarantine page, you can remove a VM from quarantine by selecting the VM and clicking **Un-Quarantine VM**.

- **Quarantine infected files**—You can specify that infected files be quarantined.

Use the Quarantine Files page on the AntiVirus module to display a list of files that are quarantined and take action.

The Quarantine Files page lets you delete an infected file, remove it from quarantine, or fetch it to remediate it according to your own process.

- **Suspend the VM**—You can suspend the VM entirely.

Use the Quarantine Files page on the AntiVirus module to display a list of files that are quarantined and take action. See [“Understanding Quarantined VMs and Files Resulting from a vGW AntiVirus On-Access Scan” on page 148](#).

To create a custom scan that allows you to specify the files to be scanned:

1. In the Step 3 Scan Engine Configuration pane, under the On-Access file types/extensions scanning selection, select the **Custom Scan** option button.
2. Select the files to scan.



NOTE: The file types and the file locations that you specify in this pane work together to clearly identify the files to scan. For example, if you select **Scan All File Types** and **Scan Only**—for example to scan only specific locations such as c:\user\share—then all the files at that location are scanned, but only those files.

- a. Select the **Scan Archives** check box to scan all files archived in various formats.



NOTE: For improved performance, do not scan archive files.

- b. Select the types of files to scan. Select one of the following options:
 - **Scan All File Types**—Scans all types of files, delimited by the selected file location.
 - **Scan Only**—Scans only specified file types, delimited by the selected file locations. You can delete file types from the provided list to exclude them from the scan.
 - **Ignore only**—Scans all types of files except the specified types.
- c. Select the locations where the files to scan reside.

For On-Access scans, when scanning files based on location, vGW Series takes into account the drive letter of the directory. For example, given the file location C:\Program, an On-Access scan scans files only in that directory. It does not scan files in the D:\Program directory, although the directory names are the same, because it acknowledges that the drive letters are different. You must specify the drive when you specify the location of files to scan for custom On-Access scans.

- **Scan All Locations**—Scans files in all locations, delimited by the selected types of files to scan.
- **Scan only**—Scans files only at the specified location, delimited by the selected types of files to scan.
- **Ignore only**—Scans all files except those that reside at the specified locations.

Related Documentation

- [Understanding vGW Series AntiVirus on page 131](#)
- [vGW AntiVirus Configuration Overview on page 139](#)
- [Understanding and Installing the vGW Endpoint on page 149](#)
- [Understanding and Configuring the vGW Series AntiVirus Settings on page 291](#)
- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)

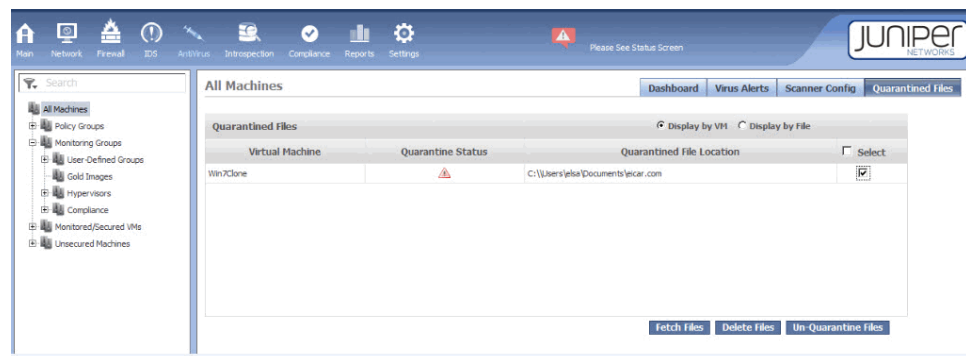
Understanding Quarantined VMs and Files Resulting from a vGW AntiVirus On-Access Scan

You use the AntiVirus module **Scanner Config** page to configure scans on groups of VMs.

A vGW AntiVirus On-Access scan can result in quarantined files or VMs:

- Quarantined files are listed in a table on the vGW AntiVirus module Quarantined Files page. Quarantining a file makes it unavailable to the VM. See [Figure 84 on page 148](#).

Figure 84: AntiVirus On-Access Quarantined Files



Use the Quarantine Files page on the AntiVirus module to display a list of files that are quarantined and take action. See . The table displays the following information for VMs with infected files: the name of the VM containing the file, the quarantine status, and the location of the quarantined file.



NOTE: You can determine how the information is displayed on the page by clicking either [Display by VM](#) or [Display by Files](#). [Display by Files](#) produces a flat, inclusive file list.

When an infected file is identified and the quarantine action is specified in the On-Access scanner configuration, the file is isolated in the vGW Endpoint on the VM. It remains there until you take action. The Quarantine Files page lets you delete an infected file, remove it from quarantine, or fetch it to remediate it according to your own process.

- Quarantined VMs are identified in the vGW Main module Quarantine page.
- Quarantining a VM effectively restricts network traffic to and from it.

Related Documentation

- [Understanding vGW Series AntiVirus on page 131](#)
- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Understanding and Installing the vGW Endpoint on page 149](#)
- [Understanding Quarantined VMs and How to Manage Them on page 157](#)
- [Understanding vGW Series on page 3](#)

Understanding and Installing the vGW Endpoint

This topic explains the vGW Endpoint and how it is used. To understand vGW Endpoint download and installation procedures within the overall context of the vGW AntiVirus configuration, see [“vGW AntiVirus Configuration Overview” on page 139](#).



WARNING: IPv4 is required for the vGW Endpoint to work properly.

- [Installing the vGW Endpoint on page 149](#)
- [vGW AntiVirus Endpoint Auto-Update on page 150](#)
- [vGW Endpoint on the VM on page 150](#)
- [Quarantined Files on page 152](#)
- [vGW Endpoint Components and Displays on page 152](#)
- [vGW Endpoint Behavior on page 153](#)

Installing the vGW Endpoint

For vGW Series On-Access scans to be performed on a VM, a vGW Endpoint must be installed on each of the VMs belonging to the VM groups specified in the On-Access scanner configuration scope. The vGW Endpoint is a binary executable (.exe file) that you can install in various ways. For example, some administrators put binaries on a network share, in which case a login script maps the drive and executes the binary. Another way to install the binary is to post it on a Web server, and download and execute it as needed. In this case, you might want to use a software package such as Microsoft Server

and Cloud Platform System Center or Manage Engine Desktop Central. You can use whatever tools you prefer for this purpose.

vGW AntiVirus Endpoint Auto-Update

After you download the vGW Endpoint and distribute it to the protected VMs in your environment that you have included in the On-Access scanner configurations, you do not need to update it. When you update the vGW Security Design VM, it automatically updates the vGW Endpoint on all VMs. That is, you install the vGW Endpoint once, and vGW Series auto-deploys an update. See [Figure 85 on page 150](#).

Figure 85: vGW AntiVirus Settings

Anti Virus Settings

☒ AntiVirus Enabled

Auto Update

☒ Enabled
☐ Disabled

AntiVirus signature update frequency (in mins): 5

Current Installed Signatures Version: Jul 29 2011, r2032

Time before declaring vGW Endpoint disconnected (in mins): 30

Time before declaring AV scan state outdated (in days): 30

Save

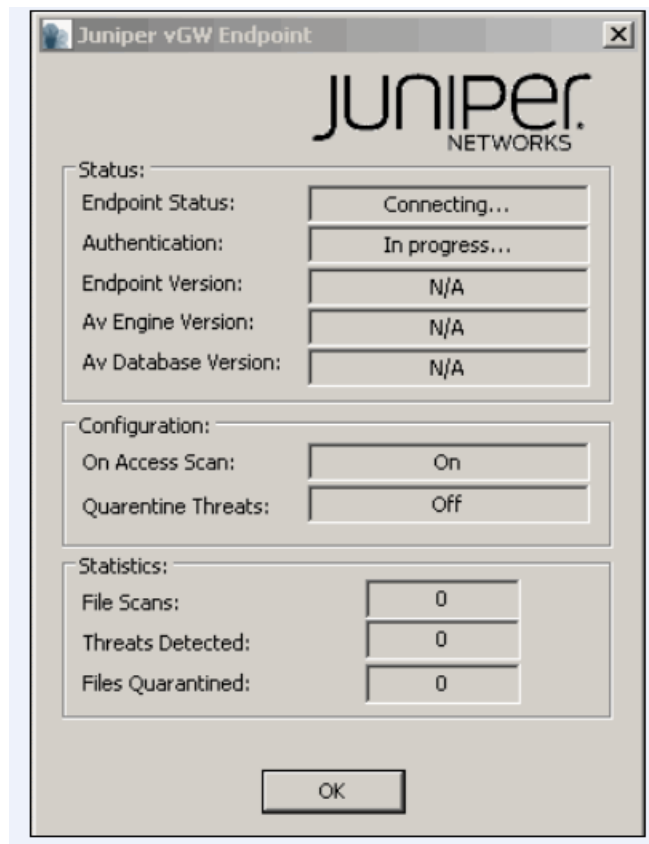
Download latest (16183.09.07.2011) vGW Endpoint installation

Download

vGW Endpoint on the VM

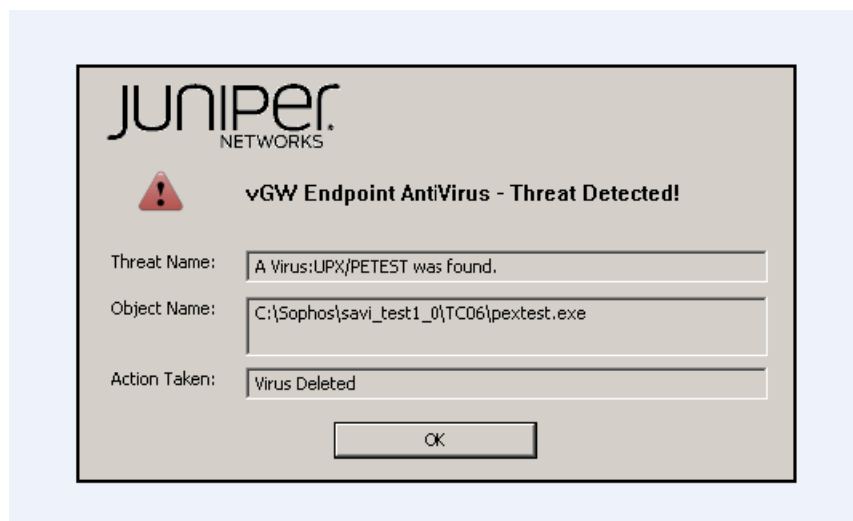
When the vGW Endpoint is connecting to the vGW Series appliance, the following dialog box appears on the VM. See [Figure 86 on page 151](#).

Figure 86: vGW AntiVirus Endpoint Connection Process Dialog Box



When vGW AntiVirus identifies a threat to the VM, it presents the following dialog box to inform you of it. See [Figure 86 on page 151](#).

Figure 87: vGW AntiVirus Endpoint Threat Detection Dialog Box



Quarantined Files

When a file is quarantined as a result of an On-Access scan, the file is sequestered in the vGW Endpoint on the protected VM. The quarantined file is inaccessible by the VM, but it remains local on it. You use the Quarantine tab on the AntiVirus module to manage quarantined files. You can handle quarantined files in these ways:

- You can fetch the file. In this case, the file is hashed and transferred off the VM for further analysis.
- You can *un-quarantine* the file. In this case, the isolated file is made available again to the VM.

In some cases, files are quarantined because of false positive results. That is, the file is suspected of being malware or infected, but that is not the case. Updating the signature database and running the scan again often resolves the problem.

- You can delete the file from the VM, if you have confirmed that it is malware or infected.



NOTE: vGW Endpoint can be used with VMware View. However, some configurations of VMware View, such as Composer, have unique configuration parameters.

For the most updated configuration information, check the JTAC Knowledge Base.

vGW Endpoint Components and Displays

The vGW Endpoint includes the following components:

- A filter driver that performs the file monitoring and scan policy enforcement.
- A service that handles communication with the vGW Security VM. It is responsible for reporting the state and enforcing the vGW AntiVirus policy, such as quarantining a file for On-Access scans.
- A tray application that reflects the known state to the service in the vGW Security Design VM. This application has three main states represented by three icons:
 - Red warning triangle—When a threat is detected, a message box appears with a red warning triangle. When the threat is dismissed, the red triangle disappears.
 - Clear burst—All components are running and connected to the vGW Security VM.
 - Burst with yellow triangle icon—The service and driver are running, but communication has not yet been fully established with the vGW Security VM.
 - Burst with red x—Either the service or the driver is not loaded. The vGW AntiVirus policy cannot be enforced in this state. When the problem is resolved, the clear burst appears.

vGW Endpoint Behavior

The vGW Endpoint captures file accesses and forwards them to the vGW Security VM for analysis. The vGW Endpoint driver caches the results of the scan. You cannot control how much of a file is transferred to the vGW Security Design VM. However, the file transfer, which is controlled internally, is efficient, based on its match against the AntiVirus signatures. Only as much of the file as is necessary to determine if it is malicious is forwarded.

The vGW Endpoint cache is not associated with a timer. For this reason, you cannot control when the cache is cleared. However, the cache is cleared during installation, un-installation, and reboot processes, when the AntiVirus signature set is updated, which is typically every few hours, and when there is a version change. You use the Settings module Security Settings > AV Settings page to specify the update frequency of the AntiVirus signatures.



NOTE: Restarting the Endpoint does not clear the cache.

Related Documentation

- [Understanding vGW Series AntiVirus on page 131](#)
- [vGW AntiVirus Configuration Overview on page 139](#)
- [Understanding and Configuring the vGW Series AntiVirus Settings on page 291](#)
- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)

Configuring vGW Series AntiVirus On-Demand Scanning

This topic explains how to configure the On-Demand vGW AntiVirus scan feature that allows you to schedule an offline full disk scan. For a smaller scan footprint, you can identify the parts of your disk that you want scanned, or you can exclude parts of it from the overall scan. To gain an overall understanding of AntiVirus configuration, before you read this topic, read "[vGW AntiVirus Configuration Overview](#)" on page 139.



NOTE: On-Demand scans are performed without any impact to the VM. The scanning is done outside the VM on the ESX/ESXi host's vGW Security VM. Therefore it not necessary to re-configure a VM after an On-Demand scan.

On-Demand scanning does not require that any software be installed in the VM. That is, you do not need to install the vGW Endpoint, which is required for On-Access scans.

On-Demand scanning can be used for many purposes. Some companies run On-Demand scans regularly to check for compliance. Public clouds that host many customer VMs

but that do not have jurisdiction to install vGW Endpoints on the VMs use On-Demand AntiVirus scanning.



NOTE: You can configure both On-Access Scanning and On-Demand Scanning in a single AntiVirus configuration.

The On-Demand scanner performs rootkit detection. The vGW AntiVirus engine contains signatures that help to identify rootkit files. It can detect files with suspicious names such as `mal.exe`, and `simpletroj.exe`. Because you do not need to protect VMs with the vGW firewall and you do not need to install the vGW Endpoint on the VM, On-Demand scans can be performed on virtual disk files from a protected location that is not compromised. This advantage increases the ability of vGW AntiVirus to detect and locate rootkits. The vGW AntiVirus engine contains signatures that help to identify rootkit files. It can detect files with suspicious names such as `mal.exe`, `simpletroj.exe`, and so on.

vGW Series scans one VM at a time to avoid problems such as brown-outs that could ensue during an On-Demand full disk scan if all VMs were scanned concurrently. The entire disk is scanned according to the schedule configuration specifications, but VMs are scanned sequentially. This approach applies also to custom scans in which only selected areas of a disk are scanned.

For On-Demand scans, vGW Series scans 500 MB per second. To gain an understanding of how long a disk scan takes, consider the following equation:

$$<VM\ memory\ size> \times <number\ of\ VMs\ on\ disk> / 500\ MB\ per\ second$$

To create an On-Demand vGW AntiVirus configuration or add a new one:

1. Select the vGW AntiVirus module. On the main vGW AntiVirus page, select the **Scanner Config** tab, and click **Add**. [Figure 88 on page 155](#) shows the configuration page that appears.

Figure 88: Scanner Config Tab

Dashboard Virus Alerts **Scanner Config**

Name:
 Description:
 Scope: [Edit](#)

Step 1 - Scan Options

☐ On-Access Scanning ☒ On-Demand Scanning

Step 2 - Scan Schedule

On-Demand Scan
☒ Daily at ☐ Weekly on at ☐ Monthly on at

Step 3 - Scan Engine Configuration

On-Access file types/extensions scanning selection:
☒ Typical Scan ☐ Custom Scan

On-Demand file types/extensions scanning selection:
☒ Typical Scan ☐ Custom Scan

Step 4 - Action

☒ Alert when a virus is detected ☒ Quarantine VM ☒ Quarantine infected files ☐ Suspend VM

[Save](#) [Cancel](#)

2. Specify a name for the vGW AntiVirus On-Demand configuration scan.
3. Select the **On-Demand Scanning** option button.
4. (Optional) Give a brief description of the configuration so that it is quickly recognizable.
5. From the All Groups list in the **Scope** box, identify the VM groups to be scanned. See [Figure 89 on page 155](#).

Figure 89: Step 2: Scan Schedule

Step 2 - Scan Schedule

On-Demand Scan
☒ Daily at ☐ Weekly on at

Step 3 - Scan Engine Configuration

On-Access file types/extensions scanning selection:
☒ Typical Scan ☐ Custom Scan

On-Demand file types/extensions scanning selection:
☒ Typical Scan ☐ Custom Scan

All Groups

- Corp-IT-Policy
- Database Policy
- Firewallled_compliant
- Hypervisors
- Linux-Corp-Policy
- Partner-DMZ (SRX - DMZ)
- PartnerDMZ
- Poomima-Test Group
- Server-AV-Group
- trust (SRX - DMZ)

Selected Groups

[Apply](#) [Cancel](#)

- In the Step 2 Scan Schedule pane, specify when you want the vGW Series to perform the scan.

You can schedule daily, weekly, or monthly scans.

- In the Step 3 Scan Engine Configuration pane, select the type of scan to perform, either Typical Scan or Custom Scan. For this example, select the **Typical Scan** option button.

Step 3 - Scan Engine Configuration

On-Access file types/extensions scanning selection:

☒ Typical Scan ☐ Custom Scan

On-Demand file types/extensions scanning selection:

☐ Typical Scan ☐ Custom Scan

Step 4 - Action

☒ Alert when a virus is detected ☒ Quarantine VM ☒ Quarantine infected files ☐ Suspend VM

Save Cancel

- In the Step 4 Action pane, specify the action to take when the scan detects a virus:



CAUTION: For On-Demand scans, you cannot quarantine files or VMs.

- Alert when a virus is detected**—The Virus Alerts tab displays information on the VMs or files that are infected.
- Suspend the VM**—You can suspend the VM entirely.

To create a custom scan that allows you to specify the files to be scanned:

- In the Step 3 Scan Engine Configuration pane, under the On-Demand file types/extensions scanning selection, select the **Custom Scan** option button.

Step 3 - Scan Engine Configuration

On-Access file types/extensions scanning selection:

☐ Typical Scan ☒ Custom Scan

☐ Scan Archives (zip,tar,tgz etc...)

☐ Scan All File Types ☒ Scan Only 386,3GR,add,ani,asp,:

☐ Ignore only bat

☒ Scan All File Locations ☐ Scan only c:\user\share

☐ Ignore only c:\temp

On-Demand file types/extensions scanning selection:

☐ Typical Scan ☐ Custom Scan

- Select the files to scan.

The file types and the file locations that you specify in this section work together to clearly identify the files to scan. For example, if you select **Scan All File Types** and **Scan Only** (specified locations, for example c:\user\share), then all the files at that location are scanned, but only those files.

Take into account the following characteristics when you configure a custom On-Demand scan:

- vGW Series recognizes the global wildcards * and ?.

For example, you could specify C:\Program Files\MS*. You could also use the wildcard on an extension, for example doc*.

- For file locations, drive letters are ignored. For example, C:\Program Files matches the following directories, and files in both these locations are scanned:

C:\Program Files and D:\Program

vGW Series performs an On-Demand scan offline and does not take into account drive letters.

Select the **Scan Archives** check box to scan all files archived in various formats. For improved performance, do not scan archive files.

3. Select the types of files to scan. Select one of the following:

- **Scan All File Types**—Scans all types of files, delimited by the selected file location.
- **Scan Only**—Scans only specified file types, delimited by the selected file location. You can delete file types from the provided list to exclude them from the scan.
- **Ignore only**—Scans all types of files except the specified types.

4. Select the locations where the files to scan reside.

- **Scan All Locations**—Scans files in all locations, delimited by the selected types of files to scan.
- **Scan only**—Scans files only at the specified location, delimited by the selected types of files to scan.
- **Ignore only**—Scans all files except those that reside at the specified locations.

Related Documentation

- [Understanding vGW Series AntiVirus on page 131](#)
- [Understanding and Installing the vGW Endpoint on page 149](#)
- [Understanding and Configuring the vGW Series AntiVirus Settings on page 291](#)
- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)

Understanding Quarantined VMs and How to Manage Them

This topic covers aspects of the vGW Series quarantine feature. When a VM is quarantined as a result of a vGW AntiVirus, Compliance, or Image Enforcer scan, the VM is added to the Quarantine Policy group in the VM tree.

When a VM is added to the Quarantine Policy group, the quarantine policy that you configured using the Firewall module is applied to it. After a VM is quarantined, at any time, you can use the Main module Quarantine tab to manage it in various ways.


The Quarantine Policy group, the quarantine policy associated with it, and the Main module Quarantine tab cooperate to help you control and manage quarantined VMs. This topic includes the following sections:

- [About vGW Series Quarantine on page 158](#)
- [Configuring a Quarantine Policy on page 159](#)
- [Viewing the Quarantined VMs, Releasing Them From Quarantine, and Resolving Problems on page 160](#)

About vGW Series Quarantine

The Quarantine Policy group belongs to the Policy Groups branch. [Figure 90 on page 159](#) shows that one quarantined VM has been added to the Quarantine Policy group.

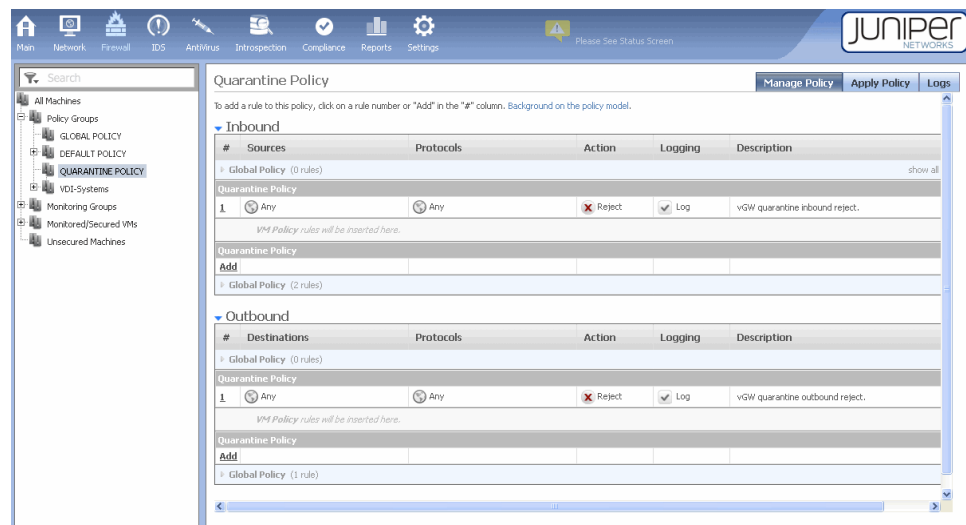
Figure 90: Quarantine Policy in the VM Tree



Configuring a Quarantine Policy

The Firewall module allows you to configure policy rules, including configuring a quarantine policy. You use the Quarantine Policy page for this purpose.

Figure 91: Configuring a vGW Series Quarantine Policy



To display the Quarantine Policy page:

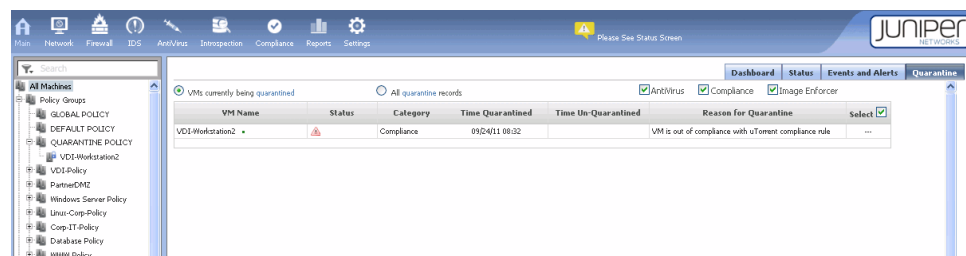
1. Select the Firewall module on the taskbar.
2. Select the Quarantine Policy group.
3. Configure the policy rules. For details on configuring policy rules, see [“Understanding the vGW Series Firewall Module” on page 93.](#)

Viewing the Quarantined VMs, Releasing Them From Quarantine, and Resolving Problems

The Main module Quarantine tab page displays a table that includes a row for each quarantined VM. You can display information for VMs quarantined as a result of vGW AntiVirus, Compliance, and Image Enforcer scans. You can display information for all quarantined VMs or VMs by scan category.

The table identifies the time the VM was quarantined and the reason for it. See [Figure 92 on page 160.](#)

Figure 92: Main Module Quarantine Tab



To view a quarantined VM in the quarantine table, resolve the problem, and remove it from quarantine:

1. Select the Main module in the taskbar.
2. Select the Quarantine tab.
3. To remove the VM from quarantine, select the VM and click **Un-Quarantine VM**.
4. Resolve the problem that caused the VM to be quarantined.

Removing a VM from quarantine does not fix the underlying problem that caused the VM to be quarantined. A VM might be quarantined because of a compliance, image enforcer, or vGW AntiVirus violation.

You can fetch the VM to resolve it offline or you can delete the VM.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series AntiVirus on page 131](#)
- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)
- [Understanding the vGW Series Enforcer Profiles Tab on page 172](#)

CHAPTER 11

vGW Series Introspection Module

This chapter includes the following topics:

- Understanding the vGW Series Introspection Module on page 163
- Understanding the vGW Series Introspection Applications Tab on page 165
- Understanding the vGW Series Introspection VMs Tab on page 168
- Understanding the vGW Series Introspection Image Enforcer Feature on page 169
- Understanding the vGW Series Image Enforcer Tab on page 171
- Understanding the vGW Series Enforcer Profiles Tab on page 172
- Understanding the vGW Series Introspection Scheduling Feature on page 176
- Understanding the vGW Series Introspection Scan Status on page 177
- Understanding the vGW Series Introspection Registry Check Feature on page 178
- Configuring the vGW Series Introspection Registry Feature on page 180

Understanding the vGW Series Introspection Module

The vGW Security Design VM Introspection module lets you monitor the software installed on guest virtual machines (VMs) in your virtual infrastructure. You can monitor software that is installed on all MS Windows VMs and some Linux VMs that support the RPM package manager when the system scans for installed applications. Without installing endpoint software in the guest VMs, vGW Series can determine which applications are installed, the operating system type (for example, for MS Windows, XP, 2003, and so on), and it can identify registry values and any applied updates (hotfixes).



NOTE: Because not all Linux VMs support RPM, we recommend that you refer to the Juniper JTAC Knowledge Base for the most current information.

When the system scans for installed applications on MS Windows VMs, it also scans registry information. Mostly the vGW Security VM performs the scans.

For Introspection, the vGW Series centralizes the scanning engine to limit disk IO, memory, and CPU consumption, and to distribute the load across responsible vGW Security VMs. Because vGW Security VMs are responsible for most of the scanning, scalability concerns are lessened, the process is faster, and introduction of new security risks is avoided.

Although most of the scanning is constrained to vGW Security VMs, both the vGW Security VM and the vGW Security Design VM engage in the process. That is, by default the scan is performed by the vGW Security VM, but it is possible to scan a VM on which the vGW Security VM is not installed. The scan can be performed by the vGW Security Design VM.



WARNING: TCP Port 902 must be open between the vGW Security Design VM and the ESX/ESXi hosts for Introspection to work properly if the vGW Security Design VM is performing it.

The Introspection module relies on taking a snapshot of a VM and analyzing it. This method guarantees that there is no adverse impact on the active VM during the scan. After the scan is complete, the snapshot is deleted immediately. The Introspection feature is supported in both IPv4 and IPv6 environments. vGW Series can mount disks that belong to VMs with either an IPv6 address or IPv4 address bound to them.

The scan does not use network packets to probe applications in the VM. Rather, it uses native VMware interfaces to examine the disk contents. This enables a fast and accurate scan. It takes only a few seconds for vGW Series to analyze the installed applications.

The ability to determine exactly which applications are installed allows the security policy for those VMs to be precise and dynamically applied. For example, you can analyze the VMs to determine which ones are running the Apache Web server. You can then place those VMs in a Smart Group and give it a name such as “webservers”. You can configure this Smart group with a policy that allows communication through HTTP/HTTPS.

The Introspection module makes it possible for you to assess applications that are installed in the environment that are secured and those that are required but are missing. For example, you can quickly identify VMs that do not have an vGW Endpoint, if the Endpoint is required. You can quarantine these VMs with a restrictive firewall policy.

Although the Introspection feature is not intended to replace a patch management solution, you can use its capabilities in this area to determine if certain hotfixes are missing. You can then quarantine the hosts without the required hotfixes until the patch management solution deploys the proper updates.

The vGW Security Design VM groups the introspection results by type (application, operating system, and hotfix). It provides graphical summary comparisons and detailed statistics about the installed software in table format.

The Introspection page includes the following tabs:

- Applications

For details, see [“Understanding the vGW Series Introspection Applications Tab” on page 165](#).

- VMs

For details, see [“Understanding the vGW Series Introspection VMs Tab” on page 168](#).

- Enforcer Profiles

For details, see “Understanding the vGW Series Enforcer Profiles Tab” on page 172.

- Scan Status

For details, see “Understanding the vGW Series Introspection Scan Status” on page 177.

- Scheduling

For details, see “Understanding the vGW Series Introspection Scheduling Feature” on page 176.

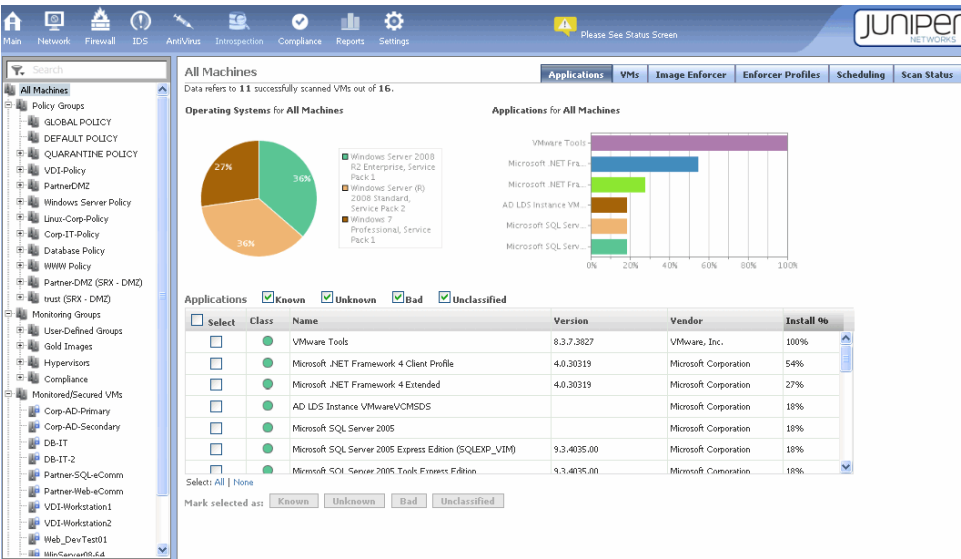
Related Documentation

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

Understanding the vGW Series Introspection Applications Tab

The Introspection module of the vGW Security Design VM includes an Applications tab that displays the following information about software currently installed on guest virtual machines (VMs). You select the VMs in the VM Tree that you want to inspect.


Figure 93: vGW Series Introspection Module Applications Tab



The Applications tab contains:


- A pie chart comparing the percentage of each type of operating system running across all secured VMs.

Operating Systems for All Machines



- A bar graph comparing the percentage of each type of application installed on all secured VMs.

Applications for All Machines



- A detailed list of each application. You can control which types of applications are included. For example, you can select only **Bad** to list applications that should not be installed on secured VMs.

Applications ☒ Known ☒ Unknown ☒ Bad ☒ Unclassified

<input type="checkbox"/> Select	Class	Name	Version	Vendor	Install %
<input type="checkbox"/>	●	VMware Tools	8.3.7.3827	VMware, Inc.	100%
<input type="checkbox"/>	●	Microsoft .NET Framework 4 Client Profile	4.0.30319	Microsoft Corporation	54%
<input type="checkbox"/>	●	Microsoft .NET Framework 4 Extended	4.0.30319	Microsoft Corporation	27%
<input type="checkbox"/>	●	AD LDS Instance VMwareVCMSDS		Microsoft Corporation	18%
<input type="checkbox"/>	●	Microsoft SQL Server 2005		Microsoft Corporation	18%
<input type="checkbox"/>	●	Microsoft SQL Server 2005 Express Edition (SQLEXP_VIM)	9.3.4035.00	Microsoft Corporation	18%
<input type="checkbox"/>	●	Microsoft SQL Server 2005 Tools Express Edition	9.3.4035.00	Microsoft Corporation	18%

Select: All | None

Mark selected as:



NOTE: If you select a group of VMs in the VM Tree, the vGW Series summarizes the data in pie and bar charts. If you select a single VM, you can view detailed information in table format.

You use the Applications tab:

- To discover information about the software installed in your environment. It provides:

- A quick overall software assessment. It allows you to quickly determine the types of installed software without regard to the exact VMs that contain it.
- The percentage of VMs running particular software. You can use this tab when you want to determine the percentage of VMs in your environment that are running a particular application, service pack, or operating system.
- Information specific to a VM or VM group. You can use this tab to discover which applications are installed on VMs or groups of VMs.
- To categorize the software installed throughout your environment. This classification system allows you to monitor the VM software state to determine if any VMs are running unauthorized or inappropriate software based on your specifications.

You can select one or more applications in the table and mark them with one of the following classifications:

- **Known**—Use this classification for applications that are acceptable for your virtualized environment.
- **Unknown**—Use this classification when an application is present, but you are unsure if it is appropriate for the environment.
- **Bad**—Use this classification for applications that are unacceptable for and prohibited from your environment.
- **Unclassified**—Use this classification when you have not yet examined an application. Newly installed applications initially show up as Unclassified.

To control displayed information:

- Click **Select All** to select all applications running in the selected VMs.
- Select **None** to clear all selected applications.
- Click a column heading in the table to sort applications by name or vendor.

The applications bar graph updates automatically as you change your selections.

Related Documentation

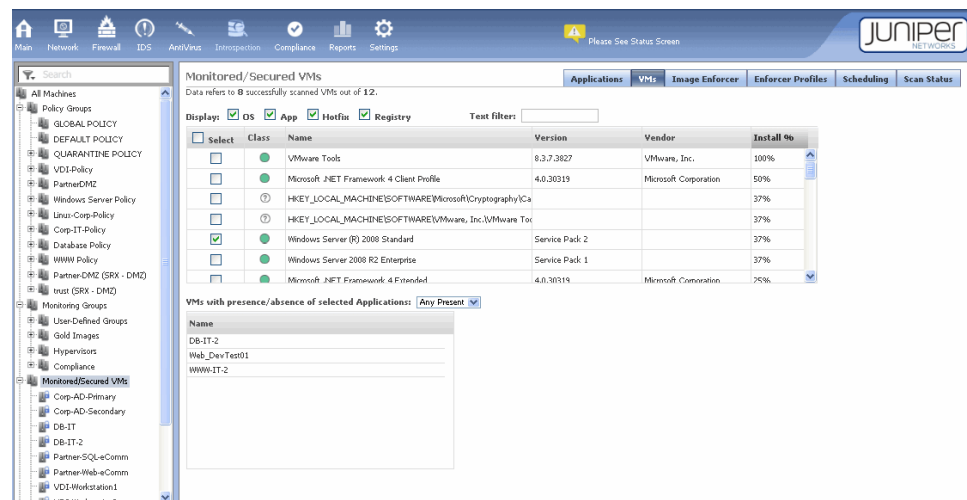
- [Understanding the vGW Series Introspection Module on page 163](#)
- [Understanding the vGW Series Introspection VMs Tab on page 168](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding the vGW Series Image Enforcer Tab on page 171](#)
- [Understanding the vGW Series Introspection Image Enforcer Feature on page 169](#)
- [Understanding vGW Series on page 3](#)

Understanding the vGW Series Introspection VMs Tab

The vGW Security Design VM Introspection feature VMs tab lets you monitor software installed on a selected VM or on a group of VMs. You can display or hide information about the operating system and about applications running on the VM, including details about installed service packs and hotfixes. You can use this feature to determine if software is present or absent on one or more VMs. The VMs tab is useful in determining which VMs have certain types of software installed.

In [Figure 94 on page 168](#), VMs in the User-Defined Groups which is selected in the VM tree are scanned to determine if they contain the Microsoft .NET Framework 4 Client Profile.

Figure 94: vGW Series Introspection Module VMs Tab



There are many ways to use this feature. For example, you can

- view all VMs that are running the MS Windows Server 2003 operating system, or all VMs that have a specific hotfix installed.
- determine the VMs that are running a specific application, such as Kazaa or Skype.
- discover VMs that are missing required software.

To search for a specific item in the list by name or vendor, click the **Name** or **Vendor** column heading in the details table, and then type the name of the software or vendor in the **Text** filter box. The list refreshes to show entries that match your specification.

You can also search the VMs to discover those that contain specific software and then filter based on a group setting in the VM Tree. To do so, select the group in the VM Tree, and then select one or more types of software in the table.

For example, select the **filter VMs with presence/absence of select Applications**, and then choose **All Present**, **Any Present**, **All Absent**, or **Any Absent** from the menu. A list of VMs meeting your criteria appears in the lower table.

vGW Series Introspection feature can discover installed software regardless of firewall settings. Because vGW Security VMs are responsible for most of the scanning, Introspection does rely on the vGW Security Design VM. That is, by default the scan is performed by the vGW Security VM. However, it is possible to scan a VM on which the vGW Security VM is not installed. In this case, the scan can be performed by the vGW Security Design VM.



WARNING: TCP Port 902 must be open between the vGW Security Design VM and the ESX/ESXi hosts for Introspection to work properly if the vGW Security Design VM is performing it.

Related Documentation

- [Understanding the vGW Series Introspection Module on page 163](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding the vGW Series Image Enforcer Tab on page 171](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Introspection Image Enforcer Feature on page 169](#)

Understanding the vGW Series Introspection Image Enforcer Feature

The vGW Security Design VM Introspection module provides a constellation of information that allows you to monitor the software installed in MS Windows and Linux guest virtual machines (VMs). It gives you deep knowledge into the state of a VM and the applications flowing between VMs, and how they are used. It can tell you the operating system versions and the services patches versions that are installed on VMs. It presents this information about the installed software to you through graphical summary comparisons and detailed statistics in table format. To facilitate management of this large amount of information and to enable you to pro-actively classify applications, the vGW Series provides an Introspection feature called the Image Enforcer.

Central to the Image Enforcer feature is the concept of a Gold Image. A Gold Image is a template from which VMs are derived, but it can also be an active VM. The Gold Image template or VM candidate has a valid and desirable configuration. When it is identified as a Gold Image, the VM is elevated to the level of a model VM configuration.

You use the Enforcer Profiles tab to create a profile for a Gold Image. In the profile, you also specify the VMs to be compared against the Gold Image and parameters that qualify the comparison. You can allow VMs to deviate from the Gold Image in various ways.

When a template is used as a Gold Image, usually the VMs that are derived from it are compared against it. For example, you might want to determine how much and in what ways their configurations have been changed since they were instantiated from the

template. However, you can specify any VMs to compare against a Gold Image, not only those that were derived from it.

You can direct the vGW Security Design VM to take certain actions based on the outcome of the comparison. For example, you can direct it to quarantine noncompliant VMs. VMs that are quarantined are viewable in the Image Enforcer page and the Main module's Quarantine page. From the Quarantine page, you can release a quarantined VM, for example, and modify it to reinstate it as a valid VM or to perform other kinds of remediation. For details on the Main module's Quarantine tab, see [“Understanding the vGW Series Main Module” on page 77](#).

You can use the Image Enforcer tab to view a summary of the comparison results and gain an overall sense of the compared VMs' conformance to the Gold Image. You can also view a bar graph specific to a particular VM to see the degree to which it conforms.

There are many ways in which to use the Image Enforcer feature:

- You might create a SQL Servers Gold Image to check for noncompliant servers.
- You might create a Desktops Gold Image and compare desktop software against it.

Consider another case. Suppose you want to use a template whose configuration is approved by auditors for PCI compliance as a Gold Image and call that Gold Image PCI-Win-Template. You could then compare the VMs belonging to the Win-PCI-Servers and PCI-Desktop VMs groups against the PCI-Win-Template Gold Image. As part of the comparison criteria, you might specify that applications classified as “known” are allowed. Although the Gold Image configuration does not contain them, a VM whose configuration contains these known applications would not be considered non-compliant.

vGW Series automatically creates a compliance rule for each Gold Image that is a template. By default, it inspects the VMs derived from the Gold Image, and it generates an alert when the compliance state changes.

You can specify when the vGW Series should scan the VMs. You can set up a scan to take place when specific events occur or based on a defined schedule that you create using the Scheduling tab. You can also limit the number of concurrent scans.

**Related
Documentation**

- [Understanding the vGW Series Introspection Module on page 163](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding the vGW Series Image Enforcer Tab on page 171](#)
- [Understanding vGW Series on page 3](#)

Understanding the vGW Series Image Enforcer Tab

The Introspection module's Image Enforcer tab reports on results of comparisons between guest virtual machines (VMs) and model templates or active VMs that are referred to as Gold Images.

Figure 95 on page 171 shows the Image Enforcer tab page displaying results of a scan in which VMs that belong to the Monitored/Secured VMs group are compared to the WinDesktop(Win7-32bit-Template) Gold Image. The groups that are included in the scan are selected in the VM tree.

Figure 95: vGW Series Introspection Module Image Enforcer Tab



The Image Enforcer tab page shows the following results of a comparison:

- It identifies matches. That is, it identifies software installed on a VM that is also installed on the Gold Image.
- It identifies applications that are installed on a VM that are not installed on the Gold Image.
- It identifies applications installed on the Gold Image that are not installed on a VM.
- It checks software versions, and it identifies versions on VMs that do not match those of the Gold Image.

Related Documentation

- [Understanding the vGW Series Introspection Module on page 163](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding vGW Series on page 3](#)

Understanding the vGW Series Enforcer Profiles Tab

This topic describes the vGW Series Introspection module's Enforcer Profiles tab. It explains how to use the Enforcer Profiles page to create profiles that allow you to compare the configurations of VMs to that of a Gold Image. It covers the information that you select or specify to create or modify a profile.

The Image Enforcer allows you to compare VMs to a VM template or an active VM that is elevated to the status of a Gold Image. For a template or an active VM to be considered a Gold Image, Gold Images are VM templates or VMs whose configurations are considered valid and desirable. Based on the outcome of the comparison scan, you can take actions such as quarantining VMs that deviate from the Gold Image, or adding or removing applications from a VM to bring it into conformance.

When VMs are quarantined, they are added to the Quarantine Policy Group. When you select a quarantined VM that is in the group, the Main module dashboard is displayed, showing compliance status for the VM, its top talkers, and IDS alerts for it. You can select the Main module Quarantine tab to take action on the VM. 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.

Before you read this topic, read [“Understanding the vGW Series Introspection Image Enforcer Feature” on page 169](#).

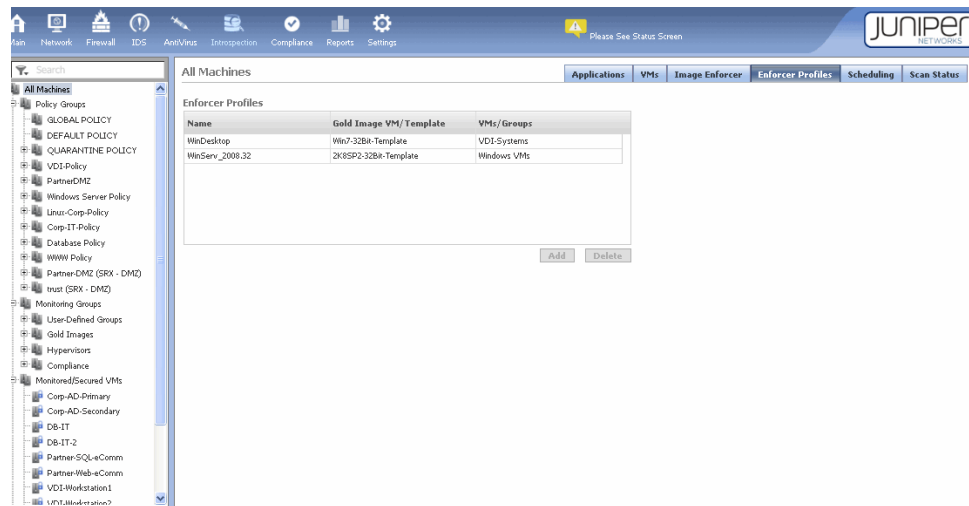
This topic includes the following sections:

- [About the Enforcer Profiles Screen on page 172](#)
- [The Add Enforcer Profile Pane on page 173](#)

About the Enforcer Profiles Screen

When you select the Introspection module Enforcer Profiles tab, the Enforcer Profiles page is displayed. Information shown in this page reflects the profiles that you have already configured, if any. You add a new Enforcer Profile from this page.

Figure 96: vGW Series Introspection Module Enforcer Profiles Tab



When you add a new profile, you give it a name that then appears in the profiles list. For each profile, the list shows the Gold Image that you selected for it and the VMs compared against it.

The Add Enforcer Profile Pane

To add a new profile, click **Add** beneath the Enforcer Profiles pane. The Add Enforcer Profile pane appears. You use this pane to configure Enforcer profiles that cover parameters for a comparison scan. In this pane, you select the Gold Image to use for the comparison; you can specify match criteria to define the comparison; and you can specify actions to take after the scan completes. You can specify conditions that exempt VMs from certain requirements, and you can specify whether the vGW Security Design VM should quarantine a non-compliant VM.

Figure 97: Adding a vGW Series Introspection Module Image Enforcer Profile

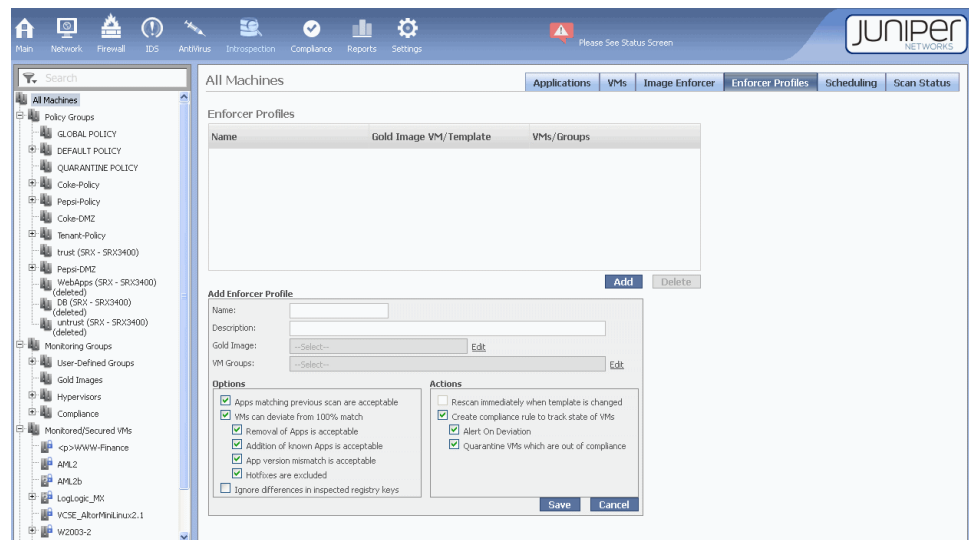


Table 9: Add Enforcer Profile: Selecting the Gold Image and VMs to Be Compared Against It

Field	Specifies
Name	A name for the profile that infers its contents.
Description	A description of the profile that indicates what it is used for.
Gold Image	<p>The VM template or VM to use as the Gold Image for this comparison. You use the Gold Image selection list to select either an existing template or VM.</p> <p>Using the option button at the bottom of the selection list, you can choose to see all Gold Image candidates or only templates or VMs.</p> <p>NOTE: After you elevate a template or VM to the status of a Gold Image, it is moved to the Gold Images group in the Monitoring Group section of the VM tree.</p>
VM Groups	<p>The VM groups or VMs whose configurations you want to compare against the selected Gold Image.</p> <p>Use the arrow buttons to include or remove a VM group or VM from the profile.</p>

Table 10: Edit Enforcer Profile Options

Option	If you select this check box, you specify that
Apps matching previous scan are acceptable	<p>If a VM was previously scanned against the profile's Gold Image and matched it, but it no longer does, the VM is allowed.</p> <p>In this case, a Gold Image might have been updated and re-scanned. Because it takes time to update the VMs specified in the Enforcer Profile group, they are allowed as matching during the transition.</p>
VMs can deviate from 100% match	A VM compared against the profile's Gold Image is allowed to deviate from it in any of the ways that you specify by selecting options identified in Table 11 on page 174 .
Ignore differences in inspected registry keys	You permit differences in registry key application settings from those of the Gold Image.

Table 11: VM Gold Image Allowed Deviations

Option	If you select this checkbox, you specify that:
Removal of apps is acceptable	An application that is missing from the VM, but that is present on the Gold Image is acceptable.
Additions of known apps is acceptable	If an application is part of a Gold Image, it is classified as known.

Table 11: VM Gold Image Allowed Deviations (*continued*)

Option	If you select this checkbox, you specify that:
App version mismatch is acceptable	The VM can contain an older or more recent version of an application than the one that exists on the Gold Image.
Hot fixes are excluded	Hot fixes are exempted from the comparison and are allowed on the VM.



CAUTION: Although you select the “App version mismatch is acceptable” option to allow a VM to contain an older or more recent version of an application than the one that exists on the Gold Image, the option might not take effect. For example, an application might have a version number as part of its program name on the MS Windows control panel. In this case, the version number might not be recognized and vGW Series would not allow the deviation. The actions that you specify in the Actions section of the Add Enforcer pane would be enacted on the VM.

Table 12 on page 175 identifies the actions that you can direct vGW Series to take following a comparison scan.

Table 12: Actions

option button	If you select this check box, you direct vGW Series to . . .
Rescan immediately when template is changed	Automatically run the comparison of the VM against the Gold Image again whenever a template that is used as a Gold Image is changed by being converted to a VM, modified, and then converted back to a template.
Create compliance rule to track state of VMs	Automatically define a compliance rule derived from the Gold Image configuration and take the actions that you select in Table 13 on page 175.

Table 13: Compliance Rule Specifications

Alert On Deviation	Notify you when the VM deviates from the Gold Image.
Quarantine VMs which are out of compliance	Quarantine VMs whose configurations do not conform with that of the Gold Image, taking into account the allowances that you specify as options described in Table 11 on page 174.

Related Documentation

- [Understanding the vGW Series Introspection Module on page 163](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)

- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding vGW Series on page 3](#)

Understanding the vGW Series Introspection Scheduling Feature

The Introspection module Scheduling page allows you to define schedules specifying when VMs are to be scanned.

To improve performance during peak periods, you can limit the number of concurrent scans by making a selection in the Max number of concurrent scans menu. We recommend running no more than two concurrent scans.

To define a scan schedule, click **Add**, select options for the scan, and then click **Save**. shows the Scheduling page with the Add Schedule dialog box displayed.

Figure 98: Introspection Module Scheduling Page

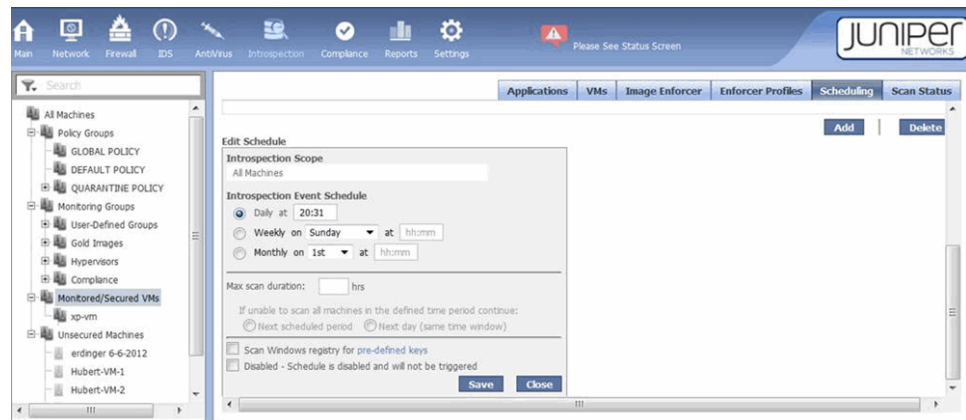


Table 14 on page 176 defines use of the fields and options.

Table 14: Scan Definition Options

Option	Select or Enter
Introspection Scope	All Machines or Selected Group , and then choose a group from the list.
Introspection Event Schedule	<p>Daily, and then enter the hour and minute when you want the scan to begin.</p> <p>Weekly, and then select the day of the week and enter the hour and minute when you want the scan to begin.</p> <p>Monthly, and then choose day of the month and enter the hour and minute you want the scan to begin.</p>
Max scan duration	The length of time that the scan must not exceed. You can use the max scan duration option to ensure that no scans occur outside maintenance. vGW Series completes a scan in progress, but it will not begin subsequent scans in the list. Any pending scans are listed in the Scan Status tab. They resume when the next scheduled time occurs.

Table 14: Scan Definition Options (*continued*)

Option	Select or Enter
If unable to scan...	Next scheduled period. The scan will continue at the next scheduled interval. Next day. The scan is continued at the same time tomorrow.
Scan Windows registry for pre-defined keys	Select the check box to direct vGW Series to inspect the registry in Microsoft OS VMs to identify user-defined registry keys and their values.
Disabled – Schedule is disabled and will not be triggered	Select the check box to disable the scan disk task in the presently defined schedule. The scan will not be performed. <i>TIP:</i> You can use this parameter to temporarily suspend scans from occurring without your having to delete the schedule then recreate it.

To delete a schedule, select the schedule in the list and click **Delete**.

Related Documentation

- [Understanding the vGW Series Introspection Module on page 163](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding vGW Series on page 3](#)

Understanding the vGW Series Introspection Scan Status

The Introspection Scan Status tab in the vGW Security Design VM lets you run and monitor disk scans of one or more VMs. vGW Series performs a full analysis of the selected VM's disk. If multiple disks exist in the VM system, each is analyzed. This analysis uncovers installed applications, the operating system, and the service pack/patch level running on the VM. You can select more than one VM in the VM tree to scan.

You can display current information about all scans (those complete and those still pending) or only complete or pending scans. You can also run scans manually or cancel scans in progress.

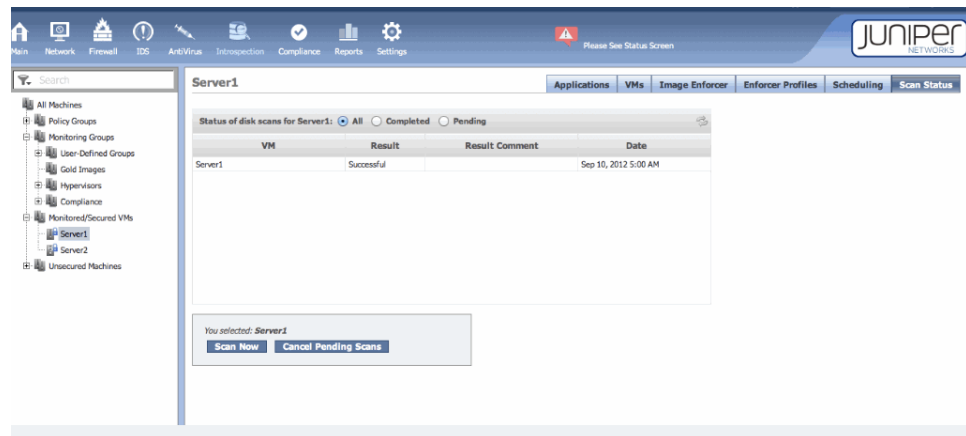
To use the Scan Status page:

- To run a scan on a selected VM or group of VMs, select the VM or VMs in the VM tree, and click **Scan Now**.
- To cancel a scan in progress, click **Cancel Pending Scans**.
- To view scan results, select **All**, **Completed**, or **Pending** to control the displayed information.

[Figure 99 on page 178](#) shows the Introspection module Scan Status page displaying the results of scans for Server1. Only one scan had been performed. The table would show

the results of all scans on Server1 if any others had been run because the All option is selected.

Figure 99: vGW Series Introspection Module Scan Status Page



vGW Series scan technology is highly accurate. Rather than a network probe, vGW Series performs an actual read of the disk file from the hypervisor. The scan process is also very fast. A typical VM scan takes less than 5 minutes.

Because scanning activity takes place on a snapshot of the system, it has no impact on the operational state of the VM. When the scan has completed, the snapshot is removed.



NOTE: vGW Series can mount disks that belong to VMs that have either IPv4 or an IPv6 address bound to them.

Related Documentation

- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Image Enforcer Feature on page 169](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)

Understanding the vGW Series Introspection Registry Check Feature

You can use the vGW Security Design VM Introspection module to inspect the registry in Microsoft OS VMs to identify user-defined registry keys and their values. You can also use it to add new registry keys.

Before you use the Settings module vGW Application Settings > Registry Values page to configure the registry introspection settings, you must be familiar with the Introspection module. For details, see [“Understanding the vGW Series Introspection Module” on page 163](#) and in particular the other topics identified in the Related Topics section of this topic.

The vGW Security VM performs the Registry inspection. It requires that the scanned VM is on a host on which the vGW Security VM resides and therefore is secured by vGW Series. However, the VM to be scanned does not need to be secured.

You can use the registry introspection feature to:

- Identify application configuration attributes. For example, it can determine if a critical directory is configured for protection by a disk encryption application.
- Validate configuration versions, such as the signature version for a DLP application in a guest VM.
- Use a registry key as internal tag to identify an MS Windows build or as an identifier for security policy automation.

You can populate the registry with values that can be used in Smart Groups for disk introspection. To configure registry introspection settings, you use the Settings module vGW Application Settings > Registry Values page. The configuration elements correlate to the registry values shown in regedit. The configuration values are:

- Name—name that identifies the registry value within vGW management.
- Key—Registry key path. Registry key names can be identified using regedit within MS Windows.
- Data—The data field contains the content associated with the chosen registry Name.



WARNING: The Key that you enter must begin with the prefix HKEY_LOCAL_MACHINE\. This is the only registry root that vGW Security VM currently supports. If the key that you enter does not contain this prefix, vGW Series displays the following alert message and highlights the Key input field.

Currently only registry values under root HKEY_LOCAL_MACHINE are supported. Please enter a key that starts with HKEY_LOCAL_MACHINE.

To add a new value:

1. Click **Add**.
2. Enter a Name for the configuration. For this example, enter **Sample Directory**.
3. Enter the registry Key. For this example, enter **HKEY_LOCAL_MACHINE\Software\Software\Sample Application\Sample Version**.
4. Enter a Value Name. For example, enter **SampleDir**.
5. Enter Data to associate with the registry key name. For example, enter **C:\Program Files\Sample Vendor\Sample**.
6. Click **Save**.

Related Documentation

- [Configuring the vGW Series Introspection Registry Feature on page 180](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)

- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Image Enforcer Feature on page 169](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Introspection Module on page 163](#)

Configuring the vGW Series Introspection Registry Feature

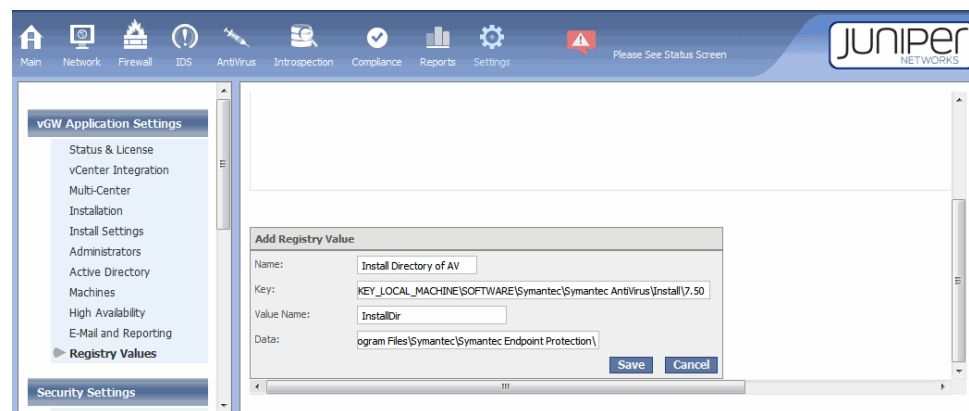
The Security Design VM Introspection > Settings > vGW Application Settings > Registry Values feature includes a disk introspection enhancement that allows you to populate a value in the registry that you can then use in Compliance inspections and in Smart Groups. The Registry Values page displays a list of registry values that are scanned on VMs during the inspection process.

Before you use the Registry Values page to configure the registry introspection settings, you must be familiar with the Introspection module. For details, see [“Understanding the vGW Series Introspection Module” on page 163](#) and in particular the other topics identified in the Related Topics section of this topic.

This example assumes that you want vGW Series to scan for data in key HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Symantec AntiVirus\Install\7.50.

1. On the Introspection > Settings > vGW Application Settings > Registry Values page, configure a new registry key. See [Figure 100 on page 180](#).

Figure 100: Configuring a New Registry Key



- a. In the **Name:** field, enter **Install Directory of AV**.
- b. In the **Key:** field, enter **HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Symantec AntiVirus\Install\7.50**.



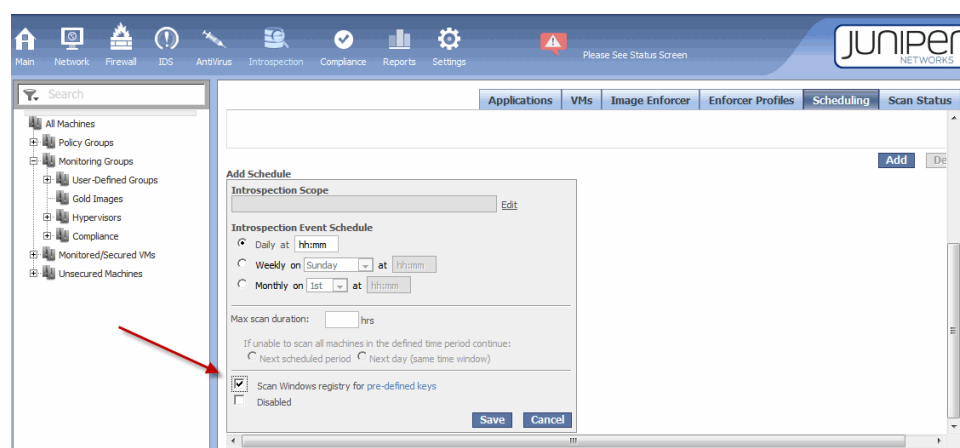
WARNING: The Key that you enter must begin with the prefix HKEY_LOCAL_MACHINE\. This is the only registry root that vGW

Security VM currently supports. If the key that you enter does not contain this prefix, vGW Series displays the following alert message and highlights the Key input field.

Currently only registry values under root HKEY_LOCAL_MACHINE are supported. Please enter a key that starts with HKEY_LOCAL_MACHINE.

- c. In the **Value Name:** field, enter **InstallDir**.
 - d. In the **Data:** field, enter **C:\Program Files\Symantec\Symantec Endpoint Protection**. This is the enforcer data.
2. To include this and all other configured Registry Values in a scheduled scan:
 - a. Check the option **Scan Windows registry for pre-defined keys** on the Introspection module Scheduling tab > Add Schedule pane. See [Figure 101 on page 181](#).

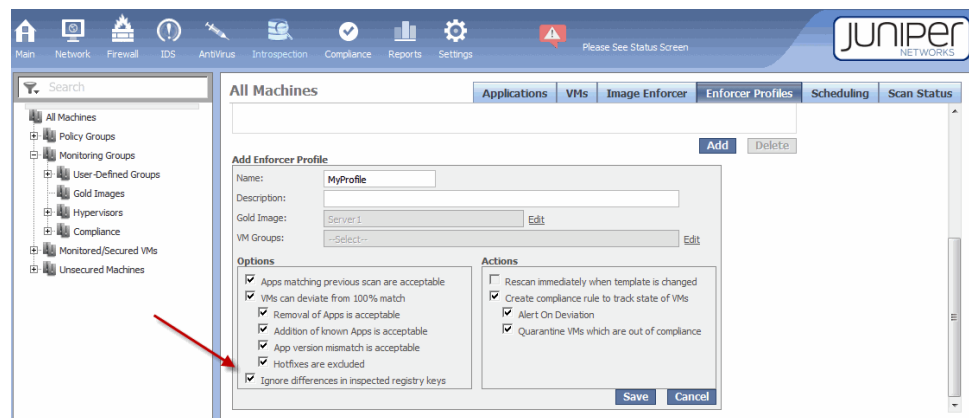
Figure 101: Add Schedule for Scan Page



To include this and all other configured Registry Values in an Enforcer Profile.

1. On the **Introspection > Enforcer Profile > Add Enforcer Profile** pane, create a profile.
2. Ensure that the **Ignore differences in inspected registry keys** check box is not selected. See [Figure 102 on page 182](#).

Figure 102: Add an Enforcer Profile that Allows for Registry Scans



Now scans that you initiate by clicking **Scan Now** on the Introspection > Scan Status page will scan registry keys.

To use registry values in a Smart Group:

1. On the Settings > Security Settings > Groups page, add your Smart Group.
2. Use the **vf.app.registry** smart property with the **contains** operator to add your condition. The value of **vf.app.registry** property will be all registry keys and their data concatenated, for example: **[key1\val1=data1,key2\val2=data2]**.

Use the Introspection module > Applications tab or the Introspection module > VMs tab to view the results of scans. The registry values will appear in the **Name** column, with their data in the **Version** column.

Related Documentation

- [Understanding the vGW Series Introspection Registry Check Feature on page 178](#)
- [Understanding the vGW Series Introspection Scheduling Feature on page 176](#)
- [Understanding the vGW Series Introspection Scan Status on page 177](#)
- [Understanding the vGW Series Introspection Image Enforcer Feature on page 169](#)
- [Understanding the vGW Series Introspection Applications Tab on page 165](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Introspection Module on page 163](#)

CHAPTER 12

vGW Series Compliance Module

This chapter covers the vGW Security Design VM Compliance module.

- [Understanding the vGW Series Compliance Module on page 183](#)
- [Configuring a Compliance Rule on page 185](#)
- [Understanding the vGW Series Hypervisor and Extended VM Security on page 188](#)

Understanding the vGW Series Compliance Module

This topic covers the Compliance module of the vGW Security Design VM that lets you monitor the compliance of your overall system with regard to industry standards best practices. Additionally, you can define rules that reflect your organization's best practices. That is, rather than using only industry best practices or standards guidelines such as PCI and HIPAA, you can define your own compliance requirements.

This topic contains the following sections:

- [The Compliance Module on page 183](#)
- [The Compliance Tab on page 184](#)
- [The Rules Tab on page 185](#)

The Compliance Module

The Compliance module relies on a rule editor that allows you to use multiple attributes about the VMware infrastructure and associated VMs to establish criteria for each designed rule. The Compliance module supports both IPv4 and IPv6 addresses. You can use any of the vGW Series built-in compliance rules in both IPv4 and IPv6 environments.

By using compliance rules to monitor key configuration parameters, you can quickly ascertain the overall state of your virtual security system. For example, you can create a compliance rule that states that non-administrative VMs are not allowed to be connected to a specific port group.

Violation of the designated rules impacts the overall compliance state. You can view details on the violations in the reports and status pages.



NOTE: If you are not using the Compliance module, you can disable it to lessen the amount of time that it takes to log into the vGW Security Design VM and improve runtime performance.

To disable the Compliance module, in `center.conf` set the following property to false:

`center.enable.compliance.onrestart`

After you change the property value, you must restart the system.

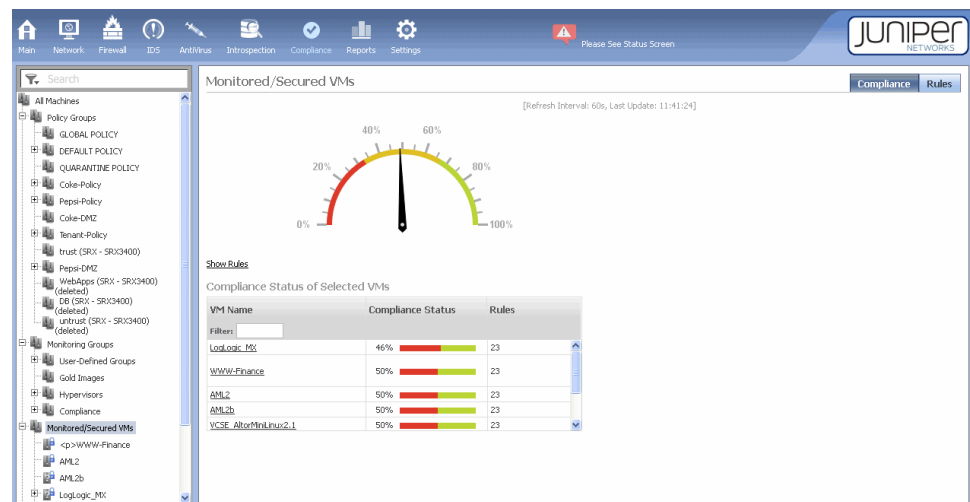
The Compliance page contains two tabs:

- Compliance
- Rules

The Compliance Tab

The Compliance tab displays a compliance meter that indicates the current level of compliance for the VM or group of VMs selected in the VM tree. It also shows statistical data that was used to calculate the overall compliance level.

Figure 103: vGW Series Compliance Module



To reflect the current compliance level, the compliance meter is refreshed automatically at 60 second intervals.

If you selected a VM group in the VM tree, the compliance meter shows the overall compliance percentage for all VMs in the group. The table below the meter lists each VM by name and shows its individual compliance level.

To display the compliance rules associated with the group, click **Show Rules**. A table appears listing each rule. It gives the name, weight, the number of VMs that the rule applies to, and the compliance status of the rule.

- To disable a rule, clear its check box.

The compliance meter is refreshed, indicating the current level of compliance with the adjusted rule set.

- Double-click a rule in the table to display details about the rule.

If you selected a single VM in the VM Tree, the compliance meter displays the current compliance of the individual machine and the rules protecting it.

The Rules Tab

The Rules tab allows you to create and manage compliance rules. This tab includes a list of defined rules that includes the name of the rule, its weight, and any labels associated with it. Labels group rules in categories.

Figure 104: vGW Series Compliance Module Rules Tab

Rule Name	Weight	Labels	Quarantine
Filters: <input type="text"/> Filter by: VMware-VM			
Backdoor Communications	1	VMware-VM	off
Clipboard enabled	1	VMware-VM	off
Disk shrink on	1	VMware-VM	off
Editable devices	1	VMware-VM	off
Forged MAC Addresses	1	DISA, NSA, VMware-VM	off
Guest MAC Address Change	1	DISA, NSA, VMware-VM	off

[Add Rule from Pre-defined List](#)
[Add](#)
[Disable](#)
[Delete](#)

You can narrow the list of rules displayed using the **Filter by** menu.



NOTE: vGW Series provides several built-in compliance rules and templates which assess the virtual infrastructure against security and hardening guidelines from VMware. These rules are also good examples to use to learn how the Compliance module works. You can use these built-in compliance rules in both IPv4 and IPv6 environments.

Related Documentation

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

Configuring a Compliance Rule

This topic explains how to create a compliance rule. For an overview of the Compliance module, see “[Understanding the vGW Series Compliance Module](#)” on page 183.

To create a compliance rule:

1. From the Compliance module > Rules tab, click Add. The Add Rule dialog box appears.

Add Rule

Name:

Comment:

Remediation:

Compliance Scope: [Edit](#)

Weight:

☐ Generate Alert when compliance state changes

☐ Quarantine non-compliant VMs

Compliance Groupings: [Edit](#)

Create Groups For:

☐ Compliant VMs

☐ Non-Compliant VMs

Advanced

Matches: ☒ All ☐ Any

[?](#) [-](#) [+](#)

[Test](#) [Save](#) [Cancel](#)

2. Define the rule. [Table 15 on page 186](#) describes the available options.

Table 15: Compliance Rule Creation Parameters

Option	Action
Compliance Scope	Select All Machines or Selected Group , and then choose a group from the list.
Name	Enter a name for the rule. Rule names can contain characters and numbers and should be descriptive, yet simple. You can describe the rule in more detail in the Comment field, if needed.
Weight	Enter a weight to be used when calculating the compliance level.
Generate Alert when compliance state changes	Direct the vGW Series to post a warning when the compliance level changes.
Compliance Groupings	Click Edit , move one or more labels to the Selected Labels list, and then click Apply .
Create Groups	<p>Create groups comprised of members who meet or violate the designated match criteria (defined in the Matches field).</p> <p>You are not required to create groups, but if you do select one of the two options, you will by default create a non-policy, Smart Group. This group can be changed to a Policy group through Settings -> Security Settings -> Groups. The benefit of automatically creating a compliance-based group is that you can easily find VMs in the VM Tree using this criterion and use the group throughout the vGW Series Table 15 on page 186.</p>

Table 15: Compliance Rule Creation Parameters (*continued*)

Option	Action
Matches	<p>Select All if the VM must meet all criteria defined in field below or Any if the VM can meet any of the criteria defined in the field below, and then choose an attribute, choose an operator, and enter a value.</p> <ul style="list-style-type: none"> To add another criterion to the rule, click +. Click - to remove a criterion from the rule.
Advanced	Enter a selection query rather than defining. For information about query syntax.

3. Click **Test**.

The vGW checks your criteria and posts a message in the Edit Rule dialog box indicating which VMs are included in the group (if any), given the criteria you specified.

4. Click **Save**.

NOTE: In addition to the items described in [Table 15 on page 186](#), you also have the option of disconnecting VMs from the network during a compliance check. By default this option is hidden because if it is used incorrectly it can cause serious problems resulting in unintended network downtime. For example, if you created a compliance rule with this action incorrectly, you could bring all VMs offline. To enable this compliance action, execute the following command from within the Web interface of the vGW Security Design VM. After it is executed, you will see a selection box called “Disconnect from the network when non compliant”.

`http:///compDisconnect?disconnect=true (or false)`

You can select a predefined rule to use. To facilitate your search for a rule, you can specify a filter.

Figure 105: Adding a Predefined Compliance Rule

Compliance Rules			
Rule Name	Weight	Labels	Quarantine
Filters: <input type="text"/>		Filter by: VMware-VM	
Backdoor Communications	1	VMware-VM	off
Clipboard enabled	1	VMware-VM	off
Disk shrink on	1	VMware-VM	off
Editable devices	1	VMware-VM	off
Forged MAC Addresses	1	DISA,NSA,VMware-VM	off
Guest MAC Address Change	1	DISA,NSA,VMware-VM	off
<div> Add Rule from Pre-defined List Add Disable Delete </div>			

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.



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.

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

Understanding the vGW Series Hypervisor and Extended VM Security

This topic covers vGW Series security for the hypervisor and VMs that aligns with VMware hardening guidelines. Before you read this topic, read [“Understanding Hypervisors and vGW Series” on page 12](#).



NOTE: To benefit from this content, you should have a general understanding of VMware hardening guidelines.

- [The Need for Hypervisor Security on page 189](#)
- [vGW Series Hypervisor and VM Security, and VMware Hardening Guidelines on page 189](#)
- [vGW Series Hypervisor and VM Security Overview on page 189](#)
- [Remediation on page 190](#)
- [Configuration Example on page 190](#)

The Need for Hypervisor Security

In full virtualization, a layer, commonly called the hypervisor or the virtual machine monitor, exists between the virtualized operating systems and the hardware. This layer multiplexes the system resources between competing operating system instances.

In the hypervisor, the virtualization infrastructure introduces a new layer of abstraction with potential exposure for malware attacks. Attempts to exploit the hypervisor as a target for attacks have increased in the recent past, and they are expected to continue to increase in number and kind in the near future. Attacks on the hypervisor can cause serious disruption such as compromise of sensitive data and denial of service (DoS). Any exposure on the hypervisor can expose guest virtual machines (VMs) that belong to many different tenants. Because the hypervisor is a crucial resource in the virtualized environment, protection of it is vital to overall security.

vGW Series enables you to verify that the hypervisor hosts that you secure meet security and compliance standards needed for a secure environment. The built-in hypervisor compliance checks are based on VMware security hardening guidelines. Additional custom hypervisor compliance checks can be created to automate any needed security compliance checks. You can use the built-in hypervisor compliance checks for hypervisors that have either an IPv4 or IPv6 address.

vGW Series Hypervisor and VM Security, and VMware Hardening Guidelines

vGW Series hypervisor security aligns with VMware hardening guidelines in all ways that are possible. vGW Series does not implement all guidelines for certain reasons. For example:

- Some guidelines, such as “NCN12 - document VLANs used on vSwitches”, pertain to behavior that can be implemented in various ways. Because of the varied implementation, the vGW Series cannot check for violations.
- Some guidelines, such as “HST02 - Ensure uniqueness of CHAP auth secret”, pertain to components inaccessible to the vGW Series. Therefore, vGW cannot perform checks on them.

vGW Series Hypervisor and VM Security Overview

You use the Security Design VM to view and configure information for the hypervisor. You can quickly view the vGW Series compliance checks that correspond with VMware

recommendations by selecting **VMware-VM** and **VMware-host** in the filter box displayed on the Compliance module Rules tab.

When you select a rule, a pane is displayed that explains the rule and the remediation action to take in response to compliance violations.

From the **Edit Rule** pane, you can modify the definition of the rule in the following ways. You can change:

- The scope of groups that the rule applies to, in the Compliance Scope list.
Click **Edit** to display the list of configured VMs and groups.
- The rule weight, in the **Weight** field, from 1–5.
- Whether an alert is generated when the compliance state of a hypervisor or a VM that belongs to the group changes.
- Whether non-compliant VMs and hypervisors should be quarantined.
- Whether the vGW Security Design VM should automatically create hypervisors groups for **Compliant VMs** and **Non-Compliant VMs**.

Remediation

For each compliance check (rule), specific remediation is suggested. You can also refer to the VMware hardening guidelines for additional information.

Configuration Example

To configure compliance requirements for hypervisors and view information about them, you use the VM Tree in conjunction with the Compliance module.

1. Under Monitoring Groups in the VM Tree, select the Hypervisors group to display the Hypervisor page.

The Hypervisors page shows the following information:

- The overall compliance status for the ESX/ESXi hosts in your virtualized environment.
 - For individual hypervisors that belong to the Hypervisors monitoring group, the Compliance Status of Selected VMs table shows the hypervisor IP address, its compliance status, and the number of compliance rules configured for it.
2. To display information about the rules configured for the hypervisors in the group, click **Show Rules**.

The Hypervisors page expands to show the following information:

- The **Compliance Rules for Selected VMs** table. This table shows the complete set of rules configured for the hypervisors. For each rule, it shows the following information:
 - The rule name.
 - The weight that is given to the rule.
 - The VMs—in this case, hypervisors—that the rule applies to.

- The **Quarantine** state, that is, whether quarantine is enabled for the rule.
 - The overall compliance status of the hypervisors that the rule is assigned to.
 - The **Compliance Status of Selected VMs** table that shows the following information:
 - The IP address of the hypervisor.
 - The compliance status of the hypervisor in regard to all the rules that are applied to it.
 - The number of rules that apply to it.
3. To display the configuration of any rule, in the Compliance Rules for Selected VMs table, click the rule name.

The **Edit Rule** pane is displayed. It shows the following information:

- Beneath the name of the rule, a **Comment** field giving a brief description of it.
- A **Remediation** field that suggests a remediation action that you can take to bring the hypervisor into conformance with the rule.

From the **Edit Rule** pane, you can modify the definition of the rule in the following ways. You can change:

- The scope of groups that the rule applies to, in the Compliance Scope list. Click **Edit** to display the list.
 - The rule weight, in the **Weight** field, from 1–5.
 - Whether an alert is generated when the compliance state of a hypervisor that belongs to the group changes.
 - Whether non-compliant hypervisors should be quarantined.
 - Whether the vGW Security Design VM should automatically create hypervisors groups for **Compliant VMs** and **Non-Compliant VMs**.
4. To customize the rule's syntax, from the **Edit Rule** pane for the rule, click **Advanced**. For details on configuring Smart Group definitions, see *Understanding vGW Series Smart Groups*.
5. Click **Test** to test the rule against hypervisors in the selected scope, after you configure the rule.
6. After you are satisfied with the rule definition, click **Save**.

Related Documentation

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

CHAPTER 13

vGW Series Reports Module

This chapter includes the following topics:

- [Understanding the vGW Series Reports Module on page 193](#)
- [Configuring a vGW Series Report on page 195](#)
- [Configuring Specifications for Automated Reports Using the vGW Series Reports Module on page 198](#)
- [Understanding vGW Series Custom Report Types on page 199](#)
- [Understanding vGW Series Network Reports on page 199](#)
- [About the vGW Series Firewall Reports on page 200](#)
- [About the vGW Series IDS Reports on page 200](#)
- [About the vGW Series Introspection Reports on page 201](#)
- [Understanding the vGW Series Compliance Report on page 201](#)
- [Understanding the vGW AntiVirus Report on page 202](#)

Understanding the vGW Series Reports Module

The vGW Security Design VM Reports module allows you to create automated reports and modify parameters for creating them, and then view the results when a report is generated.

The Reports module includes the following tabs:

- Add/Edit Reports

You use this tab to create reports. By default, the Add/Edit Reports Tab page table is empty. After you create one or more reports, they appear in the table.

- Recent Reports

The Recent Reports tab displays a table containing previously created reports. To open a report, double-click the desired report in the list. You can open it as a PDF file, or you can save it to the hard drive.



NOTE: To display a report as a PDF file, there must be a PDF viewer installed on your system.

Reports are formatted with a high-level header that includes the report name and the date the report was created.

You can create the following types of reports by selecting them in the Report Selection section when you configure a report specification:

- **Executive Summary**

The Executive Summary report provides a broad view of security and performance reports across all vGW Security Design VM modules.

- **Firewall**

The Firewall report identifies the top accepted and rejected connections processed by vGW Series firewalls.

- **Network Activity**

The Network Activity report provides a summary of network usage, including the most active VMs and top protocols observed on the virtualized network.

- **Security**

Security Report

The Security report includes:

- Top destinations for connections denied by the firewall.
- Top sources of connections denied by the firewall.
- Most common IDS alerts seen in the virtual network.
- Top sources that generated IDS alerts in the virtual network.
- Top destinations targeted by IDS alerts.

- **Introspection**

The Introspection report provides details on the applications installed on the selected VMs, and it provides a breakdown of operating systems used. It also provides an Image Enforcer report that is generated when VMs are compared to a Gold Image, which is a valid and desirable template or VM.

- **Compliance**

The Compliance report provides a detailed status view of all compliance label groupings.

In addition to these report types, you can also create custom reports.

Reports and charts displayed in reports show both IPv4 and IPv6 addresses.

**Related
Documentation**

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

Configuring a vGW Series Report

You can configure the vGW Security Design VM to produce reports on various aspects of your virtualized environment that is secured by vGW Series. You can create executive summary, firewall, network activity, security, Introspection and Compliance reports.

Reports and charts displayed in reports show both IPv4 and IPv6 addresses. In addition to current values, you can filter on IPv6 in reports.

When you add a new report, you must specify general information and the report destination and schedule information, as described in this procedure. For details on information that you configure for individual reports, see the topic for that report type.

From the vGW Security Design VM Reports module:

1. Click **Add** beneath the list of existing reports. Enter the following information in the displayed Add Report pane. See [Figure 106 on page 196](#) and

Figure 106: Adding a vGW Series Report Using the Reports Module

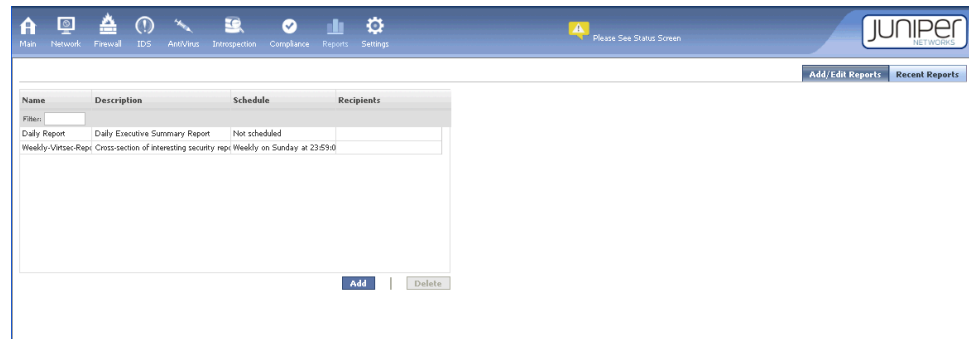



Figure 107: Defining General, Destination, and Scheduling Information for the Report



2. In the General section, specify the following information:
 - a. In the Name field, give the report a name to identify it. This is the name that is displayed in the Name column of the reports table.
 - b. Add a description that identifies the report content. This description is displayed in the Description column of the reports table.
 - c. In the Scope of report field, specify the VM groups that you want the report to cover. Click **Edit** to display a list of configured VM groups.
 - d. In the Maximum number of entries in report field, specify the record number limit.

- e. In the Time period for report field, specify the period across which data is collected. Select **Day**, **Week**, or **Month**.
 - f. In the Report format field, select the option button for either pdf or csv to specify how you want the data formatted: as a PDF file or a CSV file.
3. In the Report Destination section, specify the following information:
 - a. Whether to save the report file locally or send it to one or more recipients using e-mail.

If you select the Mail report option button, specify the source address and the addresses of recipients that you want the report mailed to.

Figure 108: Configuring the Report Destination and Generation Schedule

Report Destination

☒ Locally saved report
☐ Mail report
 From Address:
 Recipients:

Scheduling

Report Schedule:

☐ Instant report
☐ Daily at
☒ Weekly on at
☐ Monthly on at

4. Configure the report schedule.
 - To generate the report now, select the Instant report option button.
 - To generate a daily report, select the Daily at option button. In hours and minutes, specify the time when you want the report to be generated.
 - To generate a weekly report, select the Weekly on option button. Specify the day, week, and hour and minutes when you want the report to be generated.
 - To generate a monthly report, select the Monthly on option button. Specify the day of the month, and the time in hours and minutes.
5. In the Report Selection section, select the type of report to generate. You can configure the system to generate one or more types of reports. See [Figure 109 on page 197](#).

Figure 109: Configuring the Types of Reports to Generate

Administrators Filter

Username	Full Name	Type	Authentication Type
admin	Default Global Admin	Global Admin, Console administrator	Internal
admin-security-examp	admin-security-example	VM Admin	Internal

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

Configuring Specifications for Automated Reports Using the vGW Series Reports Module

This topic explains how to configure parameters that determine the kind of information to be generated in reports and when the reports should be generated.

To make report data more useful, you can use the Report Selection section to specify filters for the content. You can filter reports by Source IP, Destination IP, or Protocol. You can also filter out high, medium, and low priority alerts. Filtering allows you to report on exactly the information you need.

When you add a report specification, you can select predefined reports, including Executive Summary, Firewall, Network Activity, Security, Introspection, Compliance, and Smart Groups. You can also create custom reports.

To define a report:

1. Click **Add**.
2. Select the machines you want to report on, or select **All Machines** to report on the entire virtualized infrastructure.
3. Enter a name for the report and a description.



NOTE: Do not use spaces or special characters in the report name.

4. Specify the maximum number of entries for the report.
5. Specify the time period to report on.
6. Choose either PDF or CSV as the output format for the report.



NOTE: To display a report as a PDF file, a PDF viewer must be installed on the system.

7. Specify whether the report should be saved on the local hard disk or sent in e-mail to a recipient.



NOTE: If you use e-mail, you can specify to whom the report is sent, including an individual e-mail account, an e-mail alias, or multiple accounts separated by colons. You can also specify the e-mail address that will appear in the 'From' field on the e-mail.

8. Choose when to generate the report. You can direct the system to generate a report immediately, or you can schedule the report to run at a particular time and day.

We recommend that you schedule reports to run during low utilization periods, such as off hours. Report generation can consume significant system resources.

9. Choose a report type. Several predefined reports are provided, including Executive Summary, Firewall, Network Activity, Security, Introspection, and Compliance. You can also create custom reports.

A report selected during the report creation process has the title, graph, and relevant table data. If you select more than one type of report, each one is included in the same PDF output file, one after the other.



TIP: Select report type to display a description of the report.

10. Click **Generate Now** or **Save** to create the report.

Related Documentation

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

Understanding vGW Series Custom Report Types

This topic identifies the kinds of custom reports that you can create using the Reports module of the vGW Security Design VM. When you create a custom report, you can choose specific parameters for Network, Firewall, IDS, Introspection, and Compliance reports. Alternatively, you can use the Report Selection section of the Add Report page to select predefined reports. The core attributes of these report types are provided in the predefined reports. For an overview of the Reports module, see “[Understanding the vGW Series Reports Module](#)” on page 193.

The following topics describe the kinds of custom reports that you can define:

- [Understanding vGW Series Network Reports on page 199](#)
- [About the vGW Series Firewall Reports on page 200](#)
- [About the vGW Series IDS Reports on page 200](#)
- [About the vGW Series Introspection Reports on page 201](#)
- [Understanding the vGW Series Compliance Report on page 201](#)

Related Documentation

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

Understanding vGW Series Network Reports

This topic describes the kinds of network reports you can define using the Reports module of the vGW Security Design VM. Before you read this topic, read “[Understanding the vGW Series Reports Module](#)” on page 193.

- Top Talkers: Shows the machines generating the most traffic (combined source and destination traffic flows).
- Top Destinations: Shows where the systems are most frequently communicating.
- Top Protocols: Shows the most popular protocols in use on the virtual network.
- Top Sources: Shows which systems are generating the most traffic.
- Total Bytes: Similar to the Top Talkers report, but also shows which protocols are being used.

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

About the vGW Series Firewall Reports

vGW Series generates Firewall reports by pulling information collected by the Firewall module. You can define firewall security rules for the VMs. When connections are made to or from the resources, the firewall logs the activity and makes it available to the Reports module.

You can create the following firewall security reports:

- Top Accepted Destinations: Shows which machines in the destination field are accepting the highest number of connections, including source and destination fields for each firewall logging event.
- Top Accepted Sources: Shows the machines in the source field that are accepting the highest number of connections.
- Top Dropped or Rejected Destinations: Shows the machines in the destination field that are dropping or rejecting the highest number of connections. An action rule in a policy rule can be Accept, Drop, or Reject.
- Top Dropped or Rejected Sources: Shows the machines in the source field that are dropping or rejecting the highest number of connections.

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

About the vGW Series IDS Reports

vGW Series generates IDS reports by pulling information collected by its IDS module. These reports display a complete listing of all malicious or suspicious traffic on the virtual network.

- Top Alerts: Shows alerts seen on the virtual network
- Alert Sources: Shows sources of attacks.

The Maximum number of systems to include in the report determines how many attacks are reported. For example, if you specified a value of 20, and 20 attacks occurred on

those systems but a total of 40 attacks occurred in the specified time period, only 20 attacks would be reported.

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

About the vGW Series Introspection Reports

vGW Series generates Introspection reports by pulling information that was collected by the Introspection module. These reports contain the following information:

- **Known Applications:** Applications that you define in the Introspection module as Known. Usually they are considered good and allowed in the virtualized environment.
- **Unknown Applications:** Applications that you have determined need further investigation.
- **Bad Applications:** Applications that you have defined as bad and that are not allowed in the environment.
- **Unclassified Applications:** Applications that you have not classified. By default, a state of Unclassified indicates that the vGW Series discovered an application on a VM that it does not recognize.
- **Operating Systems:** Shows operating systems installed on VMs in the environment. vGW Series collects operating system information automatically, enabling you to run a report on all operating systems in the environment.

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

Understanding the vGW Series Compliance Report

vGW Series generates Compliance reports by pulling information collected by the Compliance module. These reports display information from the following compliance groupings:

- **DISA:** Shows information related to Defense Information Systems Agency best practices.
- **NSA:** Shows information related to National Security Agency best practices.
- **PCI:** Shows information related to Payment Card Industry best practices.
- **VMware:** Shows information related to VMware security best practices.

The resulting report shows three different summary tables containing information related to one or all of the these compliance groupings. For example, if you select just PCI and VMware, the report will contain three tables that show the values for those two compliance groupings. The first table shows all rules occurring in the selected groupings. The second table shows the groupings with summary information on rules, number of VMs, and status. The third table shows all VMs associated with the groupings.

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

Understanding the vGW AntiVirus Report

You can use the Reports module of the vGW Security Design VM to define and schedule vGW AntiVirus reports.

You can specify that the following information be included in the report:

- AntiVirus Alerts
- AntiVirus Quarantine
- AntiVirus Summary

You can specify that the report data is to be sorted by threat type, threat name, or VM.

For details on standard report configuration information that applies to AntiVirus, see [“Understanding the vGW Series Reports Module” on page 193](#).

Related Documentation

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

vGW Series Settings Module

- Understanding the vGW Series Settings Module on page 203

Understanding the vGW Series Settings Module

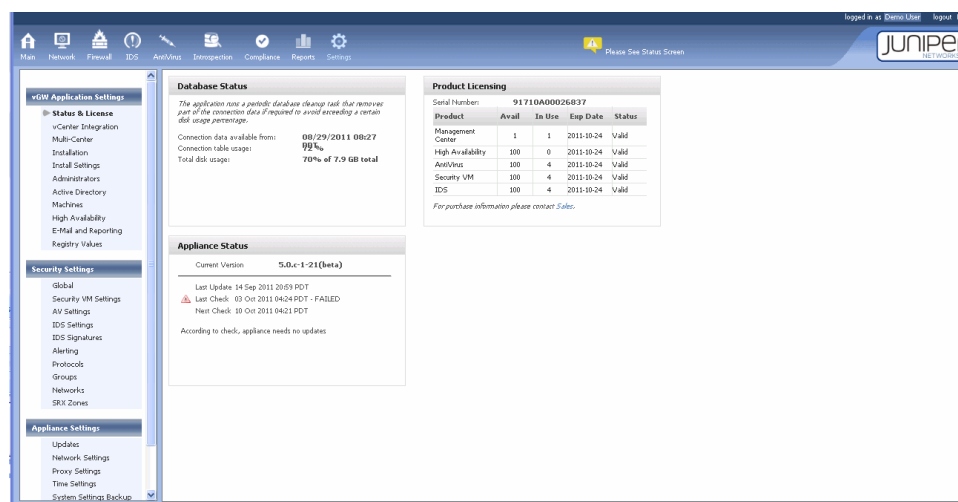
The Settings module of the vGW Security Design VM controls core vGW Series operations. The Settings module covers a wide range of information within its subsections. It contains three subsections each of which allows you to configure or view information about various parts of the system.

The Settings module contains three main sections:

- Application Settings
- Security Settings
- Appliance Settings

Figure 110 on page 203 shows the Settings module. The left navigation pane shows the sections and features that comprise the Settings module.

Figure 110: vGW Series Settings Module



- Related Documentation**
- [Understanding the vGW Series Application Settings on page 206](#)
 - [Understanding the vGW Series Security Settings on page 281](#)

CHAPTER 15

vGW Series Application Settings

- [Understanding the vGW Series Application Settings on page 206](#)
- [Viewing Status and License Information Using the vGW Series Settings Module on page 207](#)
- [Understanding Licenses for the vGW Series on page 208](#)
- [Obtaining, Installing, and Managing vGW Series Licenses on page 209](#)
- [Integrating the vGW Series with VMware Using the Settings Module on page 209](#)
- [Understanding vGW Series Integration with vCloud Director on page 214](#)
- [Configuring vGW Series Integration with vCloud Director on page 215](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Configuring vGW Series Installation Settings on page 224](#)
- [Understanding VMware Auto Deploy for ESXi Servers and with vGW Series on page 225](#)
- [Configuring VMware Auto Deploy and vGW Series to Secure ESXi Hosts on page 227](#)
- [Understanding vGW Series Timeout Parameters and the vGW Security VM Installation, Uninstallation, and Update Tasks on page 234](#)
- [Removing vGW Security VMs from ESX/ESXi Hosts on page 235](#)
- [Securing and Unsecuring Virtual Machines Using the vGW Security Design VM on page 237](#)
- [Disabling the vGW Series Suspend-Resume Process Enacted After a VM Is Unsecured on page 237](#)
- [Understanding Automatic Securing of VMs on page 238](#)
- [Understanding the vGW Series Split-Center Feature on page 240](#)
- [Understanding the Multi-Center Feature on page 244](#)
- [Configuring vGW Series Multi-Center on page 247](#)
- [Understanding vGW Series Multi-Center Synchronized Objects on page 250](#)
- [Configuring Scaling Using the Multi-Center and Split-Center Features on page 251](#)
- [Understanding the vGW Series Policy per vNIC Feature on page 258](#)
- [Configuring the vGW Series Policy per vNIC Feature on page 261](#)
- [Configuring Policy per vNIC to Secure Only Some of a VM's vNICs on page 263](#)
- [Configuring and Displaying vGW Policies for Individual vNICs on the Same VM on page 263](#)
- [Understanding Policy per vNIC and Smart Groups for VMware Environments on page 266](#)

- [Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module on page 269](#)
- [Setting Up Active Directory for vGW Series Administrator Authentication on page 275](#)
- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)
- [Configuring vGW Series E-Mail and Reporting Applications Settings on page 280](#)

Understanding the vGW Series Application Settings

The Settings module Applications section allows you to license the vGW Series product, check status on the vGW Security Design VM, control access to VMware, and add administrator information and modify it. You can also configure machines, high availability support for the vGW Security Design VM, and reporting settings.

The following topics cover specific Applications settings:

- [Viewing Status and License Information Using the vGW Series Settings Module on page 207](#)
- [Understanding Licenses for the vGW Series on page 208](#)
- [Obtaining, Installing, and Managing vGW Series Licenses on page 209](#)
- [Integrating the vGW Series with VMware Using the Settings Module on page 209](#)
- [Understanding the vGW Series Split-Center Feature on page 240](#)
- [“Understanding the Multi-Center Feature” on page 244 and “Configuring vGW Series Multi-Center” on page 247.](#)
- [Configuring Scaling Using the Multi-Center and Split-Center Features on page 251](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Configuring vGW Series Installation Settings on page 224](#)
- [Configuring the vGW Series Policy per vNIC Feature on page 261](#)
- [Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module on page 269](#)
- [Setting Up Active Directory for vGW Series Administrator Authentication on page 275](#)
- [Understanding the Multi-Center Feature on page 244](#)
- [Configuring vGW Series Multi-Center on page 247](#)
- [Configuring Scaling Using the Multi-Center and Split-Center Features on page 251](#)
- [Configuring vGW Series E-Mail and Reporting Applications Settings on page 280](#)

For details on configuring high availability for the vGW Security Design VM, see “[Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability](#)” on page 349.

Related Documentation

- [Understanding the vGW Series Settings Module on page 203](#)
- [Understanding vGW Series on page 3](#)

- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Security Design VM on page 67](#)

Viewing Status and License Information Using the vGW Series Settings Module

The Settings module Applications section allows you to view basic system status, and configure and view licensing information. It contains the following parts:

- **Database Status**—Displays the status of the internal database that stores network session data. When the database disk is full, session data for the oldest sessions is deleted. This section displays how far back session data stored in the database extends.

By default, the disk that contains the vGW Series database is set to 8 GB. If the database is not holding enough information for your environment, you can increase its size.

To increase the database size:

1. Power down the vGW Security Design VM.
2. In VMware, edit settings for the vGW Security Design VM, increasing the size of the second disk.
3. Start the vGW Security Design VM.

When the vGW Security Design VM boots up, the new disk size is recognized, and the database expands into the newly defined space.

- **Appliance Status**—This pane shows the current version of the appliance and the update version. You can use this pane to check if there are updates available that have not yet been applied.
- **Product Licensing**—This area displays a table summarizing valid licenses for the Juniper Networks vGW Series. The licensing system is 'multi-key' meaning you can attach various licenses for features and feature counts to the system.

At minimum, you must have a valid license for the vGW Security Design VM management center.

For information about licenses and specifying them, see:

- [Understanding Licenses for the vGW Series on page 208](#)
- [Obtaining, Installing, and Managing vGW Series Licenses on page 209](#)
- **Appliance Status**—This area displays the version and last update information for the vGW Security Design VM. See “System Updates” on page 76 for more information on initiating an update.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Understanding Licenses for the vGW Series

This topic contains the following sections:

- [License Requirements on page 208](#)
- [vGW Series Licenses on page 208](#)
- [Evaluation Licenses on page 209](#)

License Requirements

To enable the vGW Series, you must:

- Purchase a license for the vGW Security Design VM and separate licenses for its features, if required.
- Obtain entitlement license keys for the licenses.
- Install the vGW Series components and features license keys, as required, and manage them.

The presence of an entitlement license key determines whether you can use a feature. For information about how to purchase software licenses for vGW Series features, contact your Juniper Networks sales representative.

vGW Series Licenses

You can purchase licenses for the following vGW Series components:

- vGW Security Design VM—You must purchase a license for the vGW Security Design VM. This component serves as the management center for the vGW Series.
- Each ESX/ESXi host has a physical CPU socket count. For each host that you want to protect, you must purchase separate licenses equivalent in number to its CPU sockets. This requirement applies to the following components and features:
 - vGW Security VM—A vGW Security VM helps secure and monitor the ESX/ESXi host that it runs on, and it reports information back to the vGW Security Design VM.
 - High-availability (HA)—Allows for the deployment of primary and secondary vGW Security VMs and vGW Security Design VMs to maintain solution resiliency in the event of any single component failure.
 - AnitiVirus—Protects VMs by detecting malware, identifying affected VMs, and allowing you to define a remediation plan. The vGW AntiVirus feature does this with minimal impact to performance and resources by centralizing scanning on the vGW Security VM and, when required, using a minimal agent, called an EndPoint, on each VM.
 - Intrusion Detection System (IDS)—Allows you to examine virtual network traffic for malicious content or activity, for example, web attacks and distributed denial of service (DDOS) attacks.

Starting with two licenses, the number of licenses that you can purchase increases incrementally, depending on the license package. For Security VMs, HA, and IDS, and AntiVirus, licenses come in packages of 2, 10, 20, and upwards. You can also purchase a license for unlimited CPU sockets for each feature.

For all features except IDS and AntiVirus, licenses are perpetual. For IDS and AntiVirus, licenses are subscription-based. You can purchase a license for one year or for three years.

Evaluation Licenses

You can use an evaluation license to explore vGW Series product. The evaluation product is fully functional, and it has an embedded thirty-day license. The evaluation license enables you to use all vGW Series features. If you require a longer term license, contact your sales representative.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)
- [vGW Series Prerequisites and Resource Requirements for the VMware Environment on page 15](#)

Obtaining, Installing, and Managing vGW Series Licenses

After you power-on the vGW Security Design VM, you run the vGW Series installation wizard. During this process, you are prompted for product license information. If you purchased the vGW Series, you enter the license key for the vGW Security Design VM. After you install the license key, a serial number is presented. You can use the serial number for product support.

For each component and feature that you want to use that requires a license, you must install an entitlement license key using the vGW Security Design VM. If the proper license keys do not exist, you cannot activate the feature or install an update.

You enter license keys in the Application Settings section of the Settings module. You use the Status & License section. You install and manage licenses in the Product Licensing section. You can also use this section to view existing licenses.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Integrating the vGW Series with VMware Using the Settings Module

This topic explains the vCenter Integration settings page that allows you to configure parameters that control the interaction between vGW Series and VMware. It covers how to change the vGW Series VMware settings, direct VMware to update the vGW Security Design VM with VMs inventory information, change the settings that control how deleted VMs and information about them is handled, and how to integrate the vGW Series with the VMware infrastructure.

You can also use it to change the management domain, or scope, for the vGW Security Design VM, after you configure it initially when you install the product. The management domain specifies the data centers and host clusters in the vCenter that your vGW Security Design VM manages.

The vGW Security Design VM uses the VMware Virtual Infrastructure APIs to:

- Obtain VM Inventory information
- Determine resource utilization status
- Determine events affecting the VMs

The account used for vCenter must have read-write access to the VMware Infrastructure. You can use a custom account created in VMware; this approach makes it easier to identify and monitor activities that change. In any case, the account must have administrator privileges.

The Settings module vGW Application Settings > vCenter Integration page contains the following panes and their settings for which you either enter information or whose values you can change:

- **vCenter Settings**—Login information required for the vGW Security Design VM to communicate with the VMware vCenter and for administrator access to the vCenter. Specify the following information:
 - **Server Name or IP Address:**—Name of the vCenter or its IPv4 or IPv6 address.
 - **Username:** and **Password:**—Your administrator authentication information for accessing vCenter.
- **Scope**—Allows you to specify the vCenter's data centers and host clusters to be managed by your vGW Security Design VM. You set this value initially when you install the product. See [“Setting Up vGW Series” on page 36](#) for details on initially setting the management domain.

You use this pane to change the management domain scope. The scope for your vGW Security Design VM can be:

- **Entire vCenter**—In this case, the vGW Security Design VM is able to access and manage all VMs and other entities in all data centers in the vCenter.

To use this scope, select **Entire vCenter**.

- **Datacenter**—A subset of data centers in the vCenter.

In this case, the vGW Security Design VM is able to access and manage only the VMs and other entities in the selected data centers.

To use this scope:

1. Select **Datacenters**.

vGW Series displays all of the vCenter's data centers.



NOTE: To update the list of data centers at any time to show changes—datacenters that might have been added or removed—click **Refresh**.

2. Select the data centers for your vGW Security Design VM to manage.

Ensure that each data center is assigned to only one vGW Security Design VM. Otherwise, unexpected consequences can occur.

For an overview of the Split-Center feature, see [“Understanding the vGW Series Split-Center Feature” on page 240](#).

3. Click **Save**.

- **Clusters**—A subset of host clusters in a data center. In this case, the vGW Security Design VM is able to access only the VMs and other entities on the selected host clusters.



NOTE: All of the host clusters that you select to belong to a management domain (scope) must be in the *same* data center. You can not include host clusters from two or more different data centers in the scope.

To use the Clusters scope:

1. Click **Clusters** in the *Select a scope for your Security Design vGW* area.

In response, vGW Series displays a list of available data centers.

2. Select a data center from the displayed list whose host clusters you want the vGW Security Design VM to manage.
 - a. Click the arrow at the end of the box beside **Datacenter:** to display a list of data centers for the vCenter.
 - b. Click the data center whose cluster(s)/host(s) you want your vGW Security Design VM to manage.

vGW Series displays a list of cluster(s)/host(s) for the data center that you selected.

3. Select the check box before the names of the cluster(s)/host(s) that you want to include in your management domain.
4. Click **Save**.



NOTE: Ensure that each host cluster is assigned to only one vGW Security Design VM. Otherwise, unexpected consequences can occur.

You can change the cluster selection at any time. However, when you change the cluster scope, either of the following conditions can occur:

- Some vGW Security VMs could become unmanaged—This can occur when you remove a cluster from the list of selected clusters. Any vGW Security VM installed on an ESX/ESXi host that belongs to the removed cluster will no longer be accessible, and therefore it is no longer managed by the vGW Security VM.
- Some unmanaged vGW Security VMs could become accessible—If ESX/ESXi hosts that belong to a cluster that you add to your vGW Security Design VM management domain had a vGW Security VM installed on them by a different vGW Security Design VM, you could gain access to the vGW Security VMs. It is possible and important to gain access to an unmanaged vGW Security VM when you add its host cluster to your vGW Security VMs management domain for the following reason.

When a vGW Security VM becomes inaccessible because of cluster or datacenter selection changes its original vGW Security Design VM, its operational state might be compromised unless it is imported into another vGW Security Design vGW. This is because the vGW Security VM continues to try to communicate with its original vGW Security Design VM, which no longer recognizes it as a managed.

To view a list of unmanaged SVMs and render them manageable again:

1. Display the Settings module > Security VM Settings page.

The unmanaged vGW Security VMs are identified by a gray triangle status indicator.

2. To make a vGW Security VM manageable again, click its row to select it.
3. Click **Import**.

After you save the selection, vGW Series synchronizes all objects from vCenter. When it completes the process, vGW Series displays a message indicating the ESX/ESXi hosts and the VMs that were found.

- Deleted VMs and Groups—vGW Series can show information about any VMs and groups of VMs that it has encountered across time even if the VMs were deleted in VMware's vCenter system repository. This capability allows you to keep historic traffic records. It allows you to see all activity occurring in VMware across time. The VM's information persistency in the vGW Security Design VM can reveal attempts by a malicious administrator or hacker to bring up a VM, perform an unauthorized activity, and then delete the VM to hide their tracks.

You can change how vGW Series handles VMs that are deleted from vCenter using the following settings:

- **Hide deleted VMs from view in the Inventory Tree** check box.

By default, the "Hide deleted VMs from view in the Inventory Tree" check box is selected. However, if you do not want the deleted VMs appearing in the VM Tree, you can clear this menu item and they will be hidden from view.

The deleted VMs are still available to view again. By selecting the check box, they are again made visible in the VM Tree.

- **Delay before purging deleted VMs and Groups in days (-1 = never):** setting.

Enter the number of days after which vGW Series should purge deleted VMs and groups of VMs that have been deleted from vCenter. After that time, the VMs and all information pertaining to them is permanently deleted from vGW Series. For example, if you do not change the default value of 30 days and a VM is deleted in vCenter, at any time up to 30 days vGW Series is still able to make the VM information visible again (unhide). On the 31st day, the VM and all information pertaining to it is permanently removed from vGW Series.

- **vGW Series management server plugin**—Use this button to install the vGW Series plug-in into the vCenter interface.
 - To install the plug-in, click **Register**.
 - To view and use the plug-in, in the **vSphere Client interface** select **Home** -> **Solutions and Applications**.
 - To remove the vGW Series Management Plug-in, click **Unregister**.
- **Automatic Startup of the vGW Security Design VM and Firewall**—Use this setting to enable or disable the startup of vGW Series components when an ESX/ESXi system reboots. vGW Series components are set to start up automatically by default.
- **Synchronize machine name**—Changing the name of a VM in vCenter by default causes the name of the equivalent VM object in vGW Security Design VM to be changed to the same value. To override this setting, clear the value for this item.

For example, security administrators might want to use this override feature if they are not using the same naming convention as the VM team. The ability to override the default behavior is also useful if security administrators have created dynamic security policies using the name of the VM, and they do not want them affected by simple name changes in the vCenter.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the VMware Infrastructure and vGW Series on page 13](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [About the vGW Security Design VM Tree on page 73](#)

Understanding vGW Series Integration with vCloud Director

The vGW Security Design VM integrates directly with VMware's vCloud Director to allow vGW Series to retrieve information from vCloud Director about virtual machines (VMs). After you configure vCloud in the vGW Security Design VM, the information about a VM that it acquires can be used to dynamically associate that VM with vGW Series groups and policies that you create.

- [VMware vCloud Director on page 214](#)
- [vGW Series and vCloud on page 214](#)
- [Requirements on page 215](#)

VMware vCloud Director

VMware's vCloud Director Infrastructure-as-a-Service solution allows for rapid provisioning of complete virtual software-defined datacenter services. vCloud Director implements pooling, abstraction, and automation of data center services including storage and networking services. Using it, administrators can provision infrastructure without concern for physical hardware configuration.

Although vCloud Director can be used within an enterprise infrastructure, it is commonly used by cloud-based VM hosting providers.

vGW Series and vCloud

The vGW Security Design VM direct integration with vCloud Director allows it to collect information that is associated with a VM in vCloud Director. Information that vGW Series collects includes:

- VM membership in a specific organization.
- VM tags defined in the VM metadata. vCloud Director can associate information about VMs from its Metadata tab page that is configured by an administrator or other user, based on their permissions.

The vGW Security Design VM obtains the VM name and value data from this configuration. The vGW Security Design VM can obtain multiple values, if any.

vGW Security Design VM allows you to define Smart Groups used as policies in which VMs that match the Smart Group criteria are dynamically associated with the group, and its policy is applied to them. The vCloud Director information used in a dynamic group is associated with the `vcd.tag` property. The information appears as comma separated *attrname=value* pairs with the organization information appearing as the value for the `OrgName` attribute, such as `OrgName=Org1`.

For example, you could define a Firewall policy to be assigned to all VMs belonging to a particular organization. If the Smart Group configuration includes that organization, the Smart Group's policy is applied to the matching VM.

You might define an Introspection Image Enforcer profile that specifies that all VMs running Windows OS that belong to a particular organization must have installed on

them all applications installed on a Gold Image that they are compared to. You could also use the information acquired from vCloud Director in configuring AnitVirus scanning.

vGW Series and vCloud Director integration is characterized as follows:

- By default, vGW Security Design VM integration with vCloud Director is disabled.

To enable integration with vCloud Director, you set the `center.vcd.enabled` parameter to true:**`center.vcd.enabled=true`**.

By default it is set to false.

- vGW Series supports integration with vCloud Director 5.1 and later versions.
- Presently the vGW Security Design VM supports integration with only one vCloud Director server.

Requirements

For vGW Security Design VM to be able to integrate with vCloud Director and query it for VM inventory and other operations, the account connecting to vCloud Director must have admin privileges.

Related Documentation

- [Configuring vGW Series Integration with vCloud Director on page 215](#)
- [vGW Series Attributes for VMware on page 311](#)
- [Understanding vGW Series Groups on page 300](#)
- [Creating vGW Series Smart Groups for VMware on page 307](#)
- [Understanding vGW Series on page 3](#)

Configuring vGW Series Integration with vCloud Director

This topic covers how to integrate vGW Series with VMware's vCloud Director using the vGW Security Design VM.

Before you configure vGW Series integration with vCloud Director, you must set up vCloud Director to send relevant notifications to an Advanced Message Queuing Protocol (AMQP) broker.

vCloud Director includes an AMQP service that you can configure to work with an AMQP broker.

The Advanced Message Queuing Protocol (AMQP) is an OASIS open standard application layer protocol for message-oriented middleware. The defining features of AMQP are message orientation, queuing, routing (including point-to-point and publish-and-subscribe), reliability and security.

There are several AMQP-compatible brokers, including:

- Red Hat MRG Messaging. See <http://www.redhat.com/products/jbossenterprisemiddleware/messaging/>

- RabbitMQ. See <http://www.rabbitmq.com/>



NOTE: On the vCloud Director Administration Screen page where you configure the AMQP broker settings, you must select **Enable Notifications**. Also, set **Exchange** to **vgwExchange**. If you use a different value, ensure that it matches the value of property `center.vcd.amqp.exchange` in `center.conf`.

After you complete this configuration and you configure the vGW Security Design VM for integration with vCloud Director, the vGW Security Design VM can register with the AMQP broker to acquire these notifications and use them for updates.

To configure vGW Series integration with vCloud Director:

1. Enable vCloud Director integration by setting the `center.vcd.enabled` parameter to true. By default it is set to false.

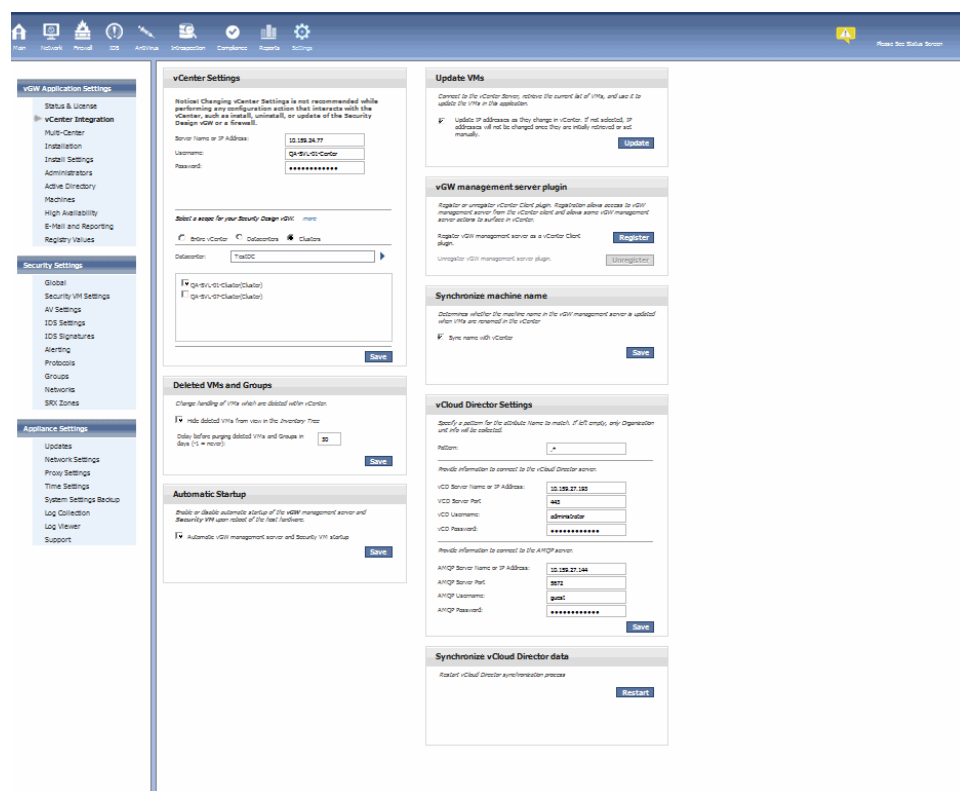
center.vcd.enabled=true

Restart Apache Tomcat.



NOTE: If you reset this value to false, all existing connections with vCloud Director are closed and the credentials are removed from the vGW Series database. In addition, the pane for configuring vCloud Director credentials is no longer displayed.

Figure 111: VGW Security Design VM vCenter Integration Window Showing vCloud Director Settings Pane



2. Select Settings > vGW Application Settings > vCenter Integration to configure vGW Series settings for integration with vCloud Director.
3. Configure the following information in the vCloud Director Settings pane:
 - Enter the IP address or DNS name of the vCloud Director server. You can specify an IPv6 or IPv4 address.
 - Specify the port number if it differs from the default of 443.

- Enter the user type. The user must have admin privileges.
- Specify a password.

4. Click **Restart** in the Synchronize vCloud Director data pane.

The vGW Security Design VM automatically configures information about any VM that it discovers through vCloud Director and it associates that information with the VM. You can view that information by selecting Settings > vGW Application Settings > Machines page.

**Related
Documentation**

- [Understanding vGW Series Integration with vCloud Director on page 214](#)
- [vGW Series Attributes for VMware on page 311](#)
- [Understanding vGW Series Groups on page 300](#)
- [Creating vGW Series Smart Groups for VMware on page 307](#)
- [Understanding vGW Series on page 3](#)

Installing vGW Security VMs on ESX/ESXi Hosts

A vGW Security VM protects and secures virtual machines (VMs) on an ESX/ESXi host where it is installed. The vGW Security VM acts as a conduit to the vGW kernel module which it inserts into the hypervisor of the host that it protects when it is installed. 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. All connections are processed and firewall security is enforced in the vGW Series kernel module. In other words, virtualized network traffic is secured and analyzed against the security policy in the vGW kernel module.

You deploy a vGW Security VM to each ESX/ESXi host in your environment that you want vGW Series to secure and monitor. The vGW Security VM protects VMs on that host and it gathers information about network traffic. It also maintains policy and logging information.

Securing an ESX/ESXi host with a vGW Security VM entails the following two parts:

- First you must install a vGW Security VM on the ESX/ESXi host to be secured. It is during this process that the vGW Security VM inserts the kernel module into the hypervisor of the ESX/ESXi host. This topic covers that process.
- Next you must select the VMs on the secured host that you want vGW Series to protect with a firewall policy and other features. The vGW Security VM obtains the policy for the VM from the vGW Security Design VM and provides the vGW Series kernel (hypervisor) module with it.

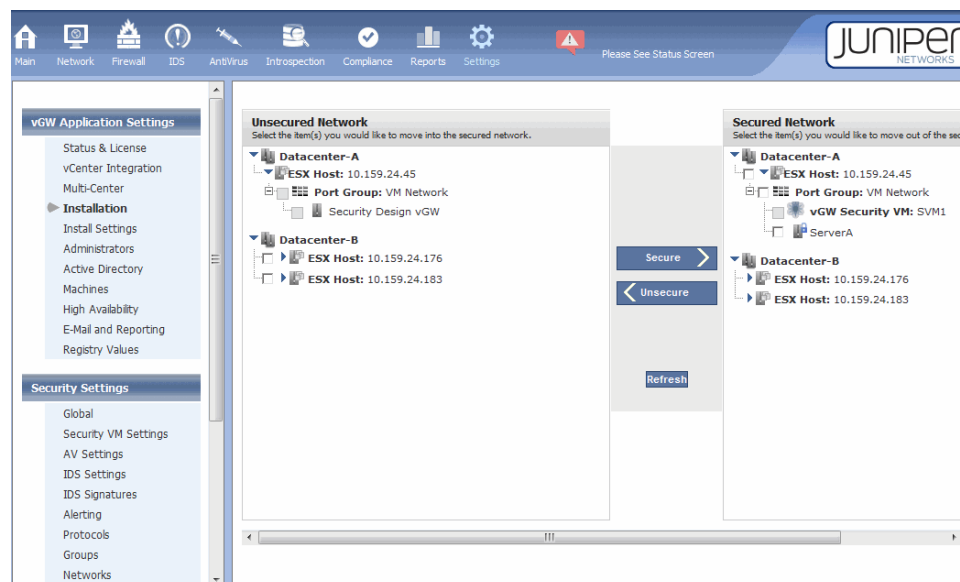
See [“Securing and Unsecuring Virtual Machines Using the vGW Security Design VM” on page 237](#) for details on the second part of the process.

To install the vGW Security VM on an ESX/ESXi host:

1. Select the Settings module **vGW Application Settings** > **Installation** page.
2. In the **Unsecured Network** pane, select the host in the data center that you want to secure with vGW Series. See [Figure 112 on page 219](#).

You can secure only one host at a time.

Figure 112: Securing an ESX/ESXi Host With a vGW Security VM

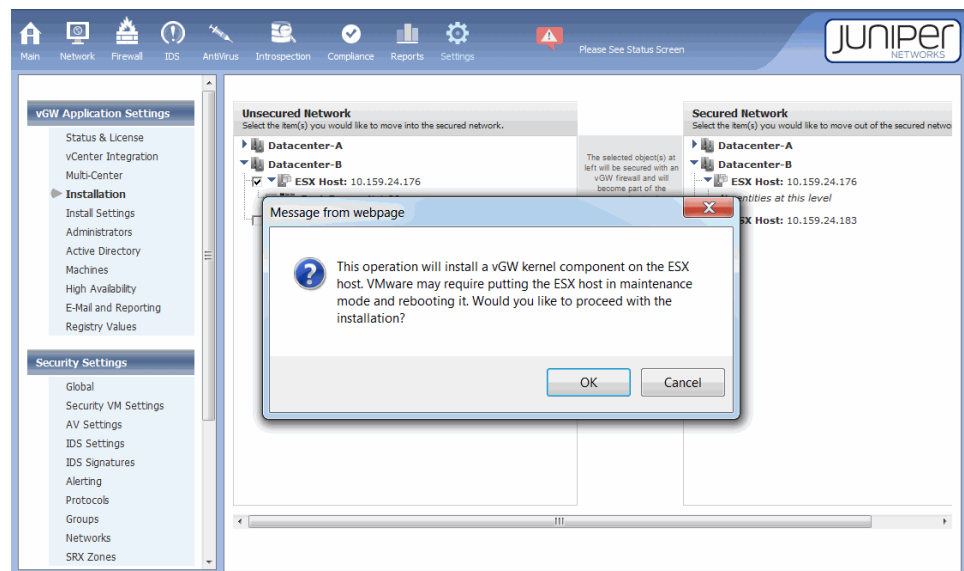


An empty check box appears before each host that is able to run the vGW Series kernel module. These hosts are not yet protected, but the check box indicates that you can secure them.

3. Click **Secure**.

After you initiate the installation process, a message is displayed indicating that VMware might require putting the ESX/ESXi host into maintenance mode and rebooting it. See [Figure 113 on page 220](#). Note that the message shown in this figure might differ somewhat depending on the vGW Series version that you are installing.

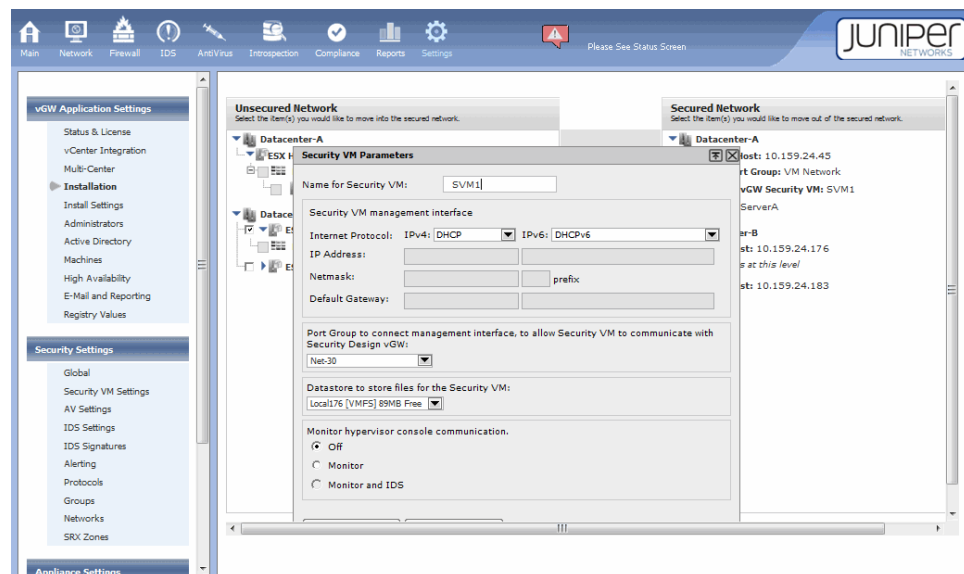
Figure 113: Installing a vGW Security VM on an ESX/ESXi Host



4. Click OK.

A dialog box is displayed allowing you to enter a name and specify other parameters for the vGW Security VM. See [Figure 114 on page 220](#).

Figure 114: Specifying vGW Security Parameters During Installation



Specify or select values for the following parameters:

- Enter a name for the vGW Security VM.
- Select the vGW Security VM security management interface addressing mode. The vGW Security Design VM communicates with the vGW Security VM

management interface based on this addressing mode. This interface must be reachable by the management interface of the vGW Security Design VM.

vGW Series supports both IPv4 and IPv6 address types. As such, the Installation Wizard for vGW Security VMs allows you to enter information for both types.

Select values for:

- IPv4
 - DHCP (Default): To obtain an IPv4 address, by default the vGW Security VM is configured to use DHCP. You do not need to specify additional information.
 - Static IP. If you select **Static IP**, you must specify a static IPv4 address and its network mask routing prefix, and the default gateway to assign to the vGW Security VM.
- IPv6
 - DHCPv6 (Default): To obtain an IPv6 address, by default the vGW Security VM is configured to use DHCPv6. You do not need to specify additional information.
 - Autoconfiguration. If you select **Autoconfiguration**, stateless address autoconfiguration is used to obtain the IPv6 address. It 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.
 - Static IP. If you select **Static IP**, you must specify a static IPv6 address, including the IPv6 address prefix (the initial bits of the address that denote the network address, akin to a netmask), and the default gateway to use for it.

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.

You can configure the vGW Security VM not to use dual stack in the following way:

- To use only IPv4 for vGW Security Design VM management communication with this vGW Security VM, disable IPv6. On the displayed list for the IPv6: box, select **Disabled**.
- To use only IPv6 for vGW Security Design VM management communication with this vGW Security VM, disable IPv4. On the displayed list for the IPv4: box, select **Disabled**.

How you configure addressing for the vGW Security VM affects its communication with the vGW Security Design VM management center. In an environment in which neither the vGW Security Design VM nor the vGW Security VM is configured for dual stack and the IP address types of their management interfaces are not the same, communication problems will occur. (For example, one interface might have an IPv6 address and the other might have an IPv4 address.) The vGW Security Design VM will not be able to connect to the vGW Security VM to carry out any procedures.

- c. Specify the port group to use to connect the vGW Security VM to the vGW Security Design VM.
- d. Specify the data store for the vGW Security VM.
- e. Specify if the hypervisor communication console should be monitored and if IDS should be used.

The dialog box allows you to enable console (hypervisor) monitoring *or* console monitoring and IDS.

- If you enable console monitoring, vGW Series monitors network traffic to the hypervisor console vNIC to ensure that inappropriate activity is not occurring.
- If you enable both console monitoring *and* IDS traffic monitoring, network traffic to the hypervisor console is monitored and IDS traffic is mirrored to the IDS engine.



WARNING: To use this option, you must first install an IDS license.

If at this point you do not enable console monitoring and IDS, you can do so later after you install a vGW Security VM. In that case, you use the Settings module Security Settings > Security VM Settings Network Monitoring tab and the IDS tab for a particular VM.

- f. Click **Secure**.

After you click **Secure**, the vGW Series associates all virtual NICs (vNICs) for the relevant VMs with the vGW Series kernel module.

VMware requires that the vNICs be disconnected and reconnected through a suspend and resume process. (VMs do not have access to the network during the few seconds that this process takes.) However, you can avoid the suspend and resume process by following the instructions covered in [“Disabling the vGW Series Suspend-Resume Process Enacted After a VM Is Unsecured” on page 237](#).

After you complete the installation, you might want to refine the configuration pertain to policy in the following ways:

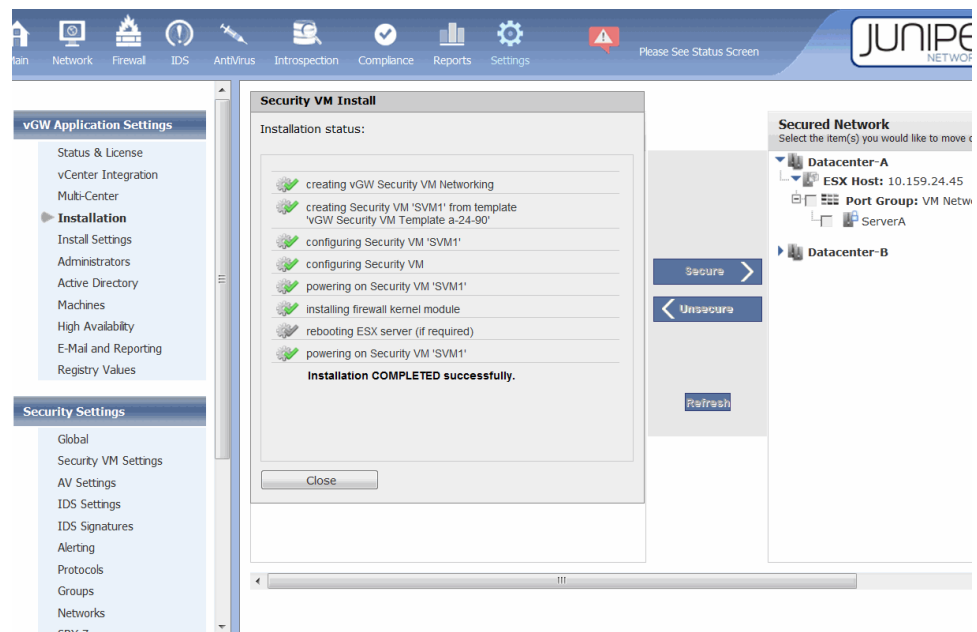
- By default, each vNIC has a restrictive default security policy. You can use the Firewall module’s Manage Policy tab to make the policy less restrictive.

- You can use the Policy per vNIC feature to configure separate firewall policies for individual vNICs on the same VM. For details on the feature, see [“Understanding the vGW Series Policy per vNIC Feature” on page 258](#) and [“Configuring the vGW Series Policy per vNIC Feature” on page 261](#).

After you define the vGW Security VM, vGW Series begins the vGW Security VM firewall installation on the selected host. It displays a progress report as it completes each task. If problems occur during the installation process, vGW Series displays messages describing them.

When the installation process is finished, vGW Series displays the list of completed tasks and the successful completion notice, as shown in [Figure 115 on page 223](#). Notice that in this case, as reported, it was not necessary to reboot the host.

Figure 115: vGW Security VM Installation Process Completion Notice



Related Documentation

- [Removing vGW Security VMs from ESX/ESXi Hosts on page 235](#)
- [Understanding the vGW Security VM on page 68](#)
- [Configuring Policy per vNIC to Secure Only Some of a VM's vNICs on page 263](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)
- [Updating vGW Security VMs in Batch Mode on page 322](#)
- [Understanding vGW Series on page 3](#)

Configuring vGW Series Installation Settings

This topic covers installation settings that you configure using the vGW Security Design VM. You use the Install Settings section of the Settings module for this purpose. The Install Settings page contains the following panes:

- VMsafe installation
- Automatic Securing of VMs
- Policy per vNIC

In the VMsafe installation pane, you can:

- Select the vGW Security VM template to use to instantiate vGW Security VMs on ESX/ESXi hosts.

From the VMsafe Template list, select the template to use.

- Specify the security behavior to follow when a vGW Security VM is unable to attach to the vGW Series VMsafe kernel module or retrieve firewall policy from the vGW Security Design VM:
 - Allow traffic to and from the vGW Security VM without security controls enforced.
 - Stop all traffic to and from the vGW Security VM. In this case, VMware disconnects the VM's vNICs.
- Specify that the vGW Security VM should only monitor the activity of the VM, but not secure it.

In this case firewall policies are not loaded onto the vGW Security VM. Monitoring mode allows you to deploy a vGW Security VM without concern that security policies will block traffic.

- Automatically secure VMs. Specify the VMs in a particular group, VMs in a policy group or with a policy applied to them, all VMs, or no VMs to be automatically secured. For details see, [“Understanding Automatic Securing of VMs” on page 238](#).

For details on installing a vGW Security VM on an ESX, see [“Installing vGW Security VMs on ESX/ESXi Hosts” on page 218](#).

If you enable the Auto-Secure feature, it automatically secures VMs and attaches security policies to them. If you choose to secure VMs automatically, you have the option of excluding a group within the selected group from being automatically secured.

For details on securing VMs or removing them from a secured network manually, see [“Securing and Unsecuring Virtual Machines Using the vGW Security Design VM” on page 237](#).

You can configure information that allows you to assign separate policies to individual vNICs.

- You use the Policy per vNIC pane to specify:
 - Whether separate policies can be configured for individual vNICs on the same VM.

- If one or more vNICs on a VM that is configured for Policy per vNIC can be exempted from having a security policy. That is, no security policy is attached to them and they are not secured by vGW Series.

You can use Policy per vNIC to apply policy rules to a vNIC that passes both IPv4 and IPv6 traffic.

For details on the Policy per vNIC feature, see [“Configuring the vGW Series Policy per vNIC Feature” on page 261](#).

For a VM with multiple vNICs, the Policy per vNIC feature allows you to use different policies for each of the vNICs. Users with VMs that connect to more than one port group/vSwitch may want different policies for each of the networks that their VMs connect to. The Policy per vNIC optional parameter, SecurePervNIC, allows you to secure some of a VM's vNICs while leaving other of its vNICs unsecured. In this case, it is the VM/port group that you secure. That is, you can use different policies for a VM based on the VM/port group. To use SecurePervNIC, you must enable Policy Per vNIC. When you use SecurePervNIC, the actual distinction is the port group, not the vNIC. That is, the vNICs of a VM are secured per VM and port group. This is due to the ambiguity of having both a secured and unsecured connection to the same Port Group. To use SecurePervNIC, you must enable Policy Per vNIC.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Security Design VM on page 67](#)

Understanding VMware Auto Deploy for ESXi Servers and with vGW Series

vGW Series allows you to secure automatically ESXi hosts generated through the VMware Auto Deploy feature. This topic covers Auto Deploy and vGW Series automatic installation of vGW Security VMs for these hosts. It includes the following sections:

- [About VMware Auto Deploy on page 225](#)
- [vGW Series Support for Auto Deploy on page 226](#)
- [vGW Series Automatic Installation of vGW Security VMs on page 226](#)

About VMware Auto Deploy

VMware Auto Deploy leverages the network Preboot Execution Environment (PXE) to rapidly provision large numbers of ESXi hosts to efficiently and easily managing their hypervisor installation and upgrades. ESXi hosts that are deployed through Auto Deploy are automatically added to a host cluster. New hosts are provisioned based on user-defined specifications. You can define specifications for various hypervisor images and host profiles to be used for different hosts.

After an ESXi host is network-booted from a central Auto Deploy server, a software image is installed on it and a vCenter host profile is then used to configure the host. When this process is done, the ESXi host is connected to vCenter, where you can create virtual machines (VMs). Apart from defining rules governing images and profiles for collective

use, this process is entirely automated, allowing for quick provisioning without user intervention.

vGW Series Support for Auto Deploy

vGW Series is designed to work in tandem with VMware Auto Deploy. It complements VMware Auto Deploy by allowing you to automatically secure ESXi hosts. You can configure JunosV Firefly to automatically install vGW Security VMs on these hosts based on clusters that they belong to, on all ESXi hosts created through auto-deploy, or on none of them, effectively disabling the feature.

vGW Series assigns a name to an automatically installed vGW Security VM based on a prefix that you specify (SVM Name prefix) when you configure vGW Series auto deploy support and an octet derived from the host's IP address.

vGW Series Automatic Installation of vGW Security VMs

vGW Series detects if an ESXi host has been added to the clusters that you selected when you configured vGW Series Auto Deploy support. For ESXi hosts in a selected cluster, it determines if a vGW Security VM is already installed on that host.

- If a vGW Security VM is already installed, vGW Series ensures that networking is set up to properly handle hosts that have been rebooted. (Restoring the network restores connectivity between the vGW Security VM and the fastpath module.)
- If a vGW Security VM is not installed on the host, vGW Series treats the ESXi host as one that was added to the cluster by the VMware Auto Deploy process. In this case, it follows the same process that it uses to install a vGW Security VM under normal conditions except for the following actions:
 - It verifies that the port group and the data store exist.
 - It omits the step that installs the fastpath module and the step that reboots the ESXi host because it is assumed the fastpath module was already embedded in the image that was deployed on the host.
- If a failure occurs, it generates an alert as shown in [Figure 116 on page 226](#).

Figure 116: vGW Series Failure Alert

Dashboard Status Events and Alerts Quarantine			
Recent Alerts and Events Consolidated Search			
Security Alerts			[Refresh Interval: 60s, Last Update: 05/14/13 10:48:42] Filter
Priority	Date	Alert	
H	05/14/13 10:45	firewall "svm_30" Kernel module status changed to ok	
H	05/14/13 10:45	firewall "svm_30" vf fpstatus changed to ok	
M	05/14/13 10:43	AutoStart option is turned off on some hosts. Host(s): 10.10.10.30, more	
H	05/14/13 10:36	Auto Deploy SVM install failed on host 10.10.10.32, details: HostCommunication	
H	05/14/13 10:35	Auto Deploy SVM install failed on host 10.10.10.31, details: HostCommunication	
H	05/14/13 10:12	firewall "svm_32" Communication status changed from never communicated to communicating	
H	05/14/13 10:12	firewall "svm_32" svm config changed to ok (Security VM configuration is ok)	
H	05/14/13 10:12	firewall "svm_32" Time synchronization status changed to time synced	
H	05/14/13 10:12	firewall "svm_32" High Availability status changed to active	

- Related Documentation**
- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
 - [Understanding the vGW Security VM on page 68](#)
 - [Understanding the vGW Security Design VM on page 67](#)

Configuring VMware Auto Deploy and vGW Series to Secure ESXi Hosts

You can configure vGW Series to monitor clusters for ESXi hosts that are provisioned through VMware Auto Deploy to install vGW Security VMs on them automatically.

This topic first explains how to set up VMware for Auto Deploy.

- [Configuring Auto Deploy in VMware on page 227](#)
- [Configuring vGW Series Auto Deploy Support on page 232](#)

Configuring Auto Deploy in VMware

VMware for Auto Deploy allows you to deploy ESXi 5.0 hosts and their associated configurations automatically.

To set up VMware Auto Deploy, you install a vCenter server, a vSphere client, the Auto Deploy service, a DHCP server, a TFTP server, and the Image Builder PowerCLI and Powershell. PowerCLI and Powershell is a commandlet and scripting language that allows you to build ESXi-based images and create rules to push out those images to your ESXi hosts.

The Auto Deploy service is a Web server that serves up ESXi images. The Auto Deploy service is embedded in the vCenter server appliance (vApp). When you install that appliance, Auto Deploy is automatically configured.

However, to configure Auto Deploy completely you specify the location for the repository where the ESXi images are stored and the repository size. You also configure an Auto Deploy connection to the vCenter server, and you specify the IP address that the Auto Deploy service should use to communicate with the network.

This topic provides information based on VMware instructions. If you encounter problems configuring VMware Auto Deploy, refer to the VMware documentation.

This process requires the following virtual machines (VMs), connectivity, and components:

- VMs. You must create VMs for:
 - VMware vCenter.
 - VMware Auto Deploy.

Install the Auto Deploy service on the same VM as vCenter, preferably.

- VMware vSphere to run PowerCLI, which requires Powershell.
- A DHCP server.
- A TFTP server.

A TFTP server is required to push the boot loader to the ESXi host. You can install it wherever you choose, including on the same VM as the vCenter server. You can install any TFTP server, for example, SolarWinds or Open TFTP.

When you install the TFTP server:

- Disable Internet Explorer ESC in MS Windows. If it is not disabled, error messages are generated reporting that you do not have access permission.
- Ensure that the timeout settings allow sufficient time to boot at least four ESXi hosts concurrently.
- Ensure that access to the TFTP server is granted.

After you complete the installation process, the TFTP folder contains the boot loader that is streamed to the ESXi host.

- There must be at least one ESXi host whose MAC address you know.
- You must have control over the IP assignment on the network.
- You must create a virtual switch (vSwitch) in VMware for vGW Series with two port groups: one for the vGW Series fastpath driver module and another for the vGW Security VM.

There are two ways to create and configure a vSwitch.

- You can use the vSphere Client Add Network wizard. It guides you through processes to create a virtual network, including how to create a vSwitch.
- You can use the VMware vSphere Configuration > Networking > Virtual Switch view for the selected ESXi host.

You must use the following names for the vSwitch and the fastpath driver port group.

- Use `vmervice-vswitch` as the name for the vSwitch.
- Use `vmervice-vmknic-pg` as the fastpath driver port group name.

There are no requirements for the name that you give to the vGW Security VM port group. Normally the vGW Security Design VM generates this name based on the vGW Security VM ID. Because the vGW Security VM does not yet exist, this information is not available.

When you configure VMware Auto Deploy, you create one or more ESXi image profiles that are used to configure the ESXi hosts that Auto Deploy generates. You can clone an existing profile using the **`new-esximageprofile`** command. In that case, specify the name of the existing profile as the value of **`-cloneprofile`**.

You must derive the image profile name from the name of the ESXi software depot. It must follow the version number of the depot file that you retrieve during the configuration process, for example **`“VMware-ESXi-5.0.0-469512-standard”`**.

If the value that you specify for **`-cloneprofile`** generates an error, for example, because the VMware naming scheme has changed, you can retrieve a list of profiles to find the correct profile name. The name should have the same six-digit build number as that of the depot ZIP file. To get a list of profile names, in PowerCLI enter:

“Get-EsxImageProfile?”

You must configure vGW Series to automatically install vGW Security VMs on the ESXi hosts provisioned by VMware Auto Deploy.

- Install the fastpath driver version that vGW Series installs normally when it installs vGW Security VMs.
- Set the net.dvfilterbindipaddress (Net.DVFilterBindIpAddress) property for the selected ESXi host to 169.254.65.1. In vSphere select **Configuration** > **Software**, and click **Advanced Settings**.

To install the Auto Deploy service and to create an Auto Deploy image profile:

1. Install vCenter server 5.0, if it is not already installed.
2. Install the vSphere 5.0 client, if it is not already installed.
3. Install the Auto Deploy service.

To verify that the Auto Deploy service is connected and configured, in vSphere, click **Home** > **vCenter Service Status**.

4. Install a TFTP server.
 - a. In vCenter, select **Home** > **Administration** > **Auto Deploy**, and then click **Download TFTP Boot Zip**.
 - b. Download the ZIP file and extract the contents to the root folder on the TFTP server.
 - c. Configure the TFTP server and start the server instance.

5. Install the PowerCLI and Powershell.

Change the execution policy.

set-executionpolicy remotesigned

6. Get the required images.

- Using PowerCLI, download the ESXi software depot (repository).

This is not the ISO image. It is a VMware file that has a name similar to the following one:

VMware-ESXi-5.0.0-469512-depot.zip

- Get the vGW Series VIB ZIP file.

There are no restrictions on where you download it.

7. Using PowerCLI, create the Auto Deploy image profile and add the ESXi image to it.

The image profile contains all modules and features that you want bundled.

All ZIP file names must include the .zip extension.

- a. Connect to vCenter.

```
connect-viserver localhost
```

- b. Add the ESXi depot.

```
add-esxsoftwaredepot ESXi-depot-zip-full-path
```

- c. Add the vGW Series VIB ZIP file.

```
add-esxsoftwaredepot VIB-zip-full-path
```

Specify the full path to where you downloaded the VIB ZIP file and include the ZIP filename.

- d. Create an ESXi image profile and add the ESXi image to it.

```
new-esximageprofile -cloneprofile "VMware-ESXi-5.0.0-469512-standard" -name  
"image-profile-name"
```

- e. Add the vGW Series VIB to the image profile.

```
add-esxsoftwarepackage -imageprofile "image-profile-name" --softwarepackage  
dvfilter-altor-vf
```

To obtain image profile software package names, enter:

```
get-esxsoftwarepackage
```

- f. Create a deploy rule. The deploy rule downloads and installs all of the modules for the image into the Auto Deploy repository.

```
new-deployrule -name "auto-deploy-rule-name?" -item "image-profile-name?"  
-AllHosts?
```

- Specify the name of the image profile that you created previously.
- Specify a name for the deploy rule to add to the image profile (*"auto-deploy-rule-name"*).



NOTE: The example rule specifies that the image applies to all hosts (-AllHosts?).

- g. Create an image profile ZIP file and export it to where you want the file to reside.

```
export-esximageprofile -imageprofile "image-profile-name?" -exporttobundle  
-filepath image-profile-location-full-pathname.
```




WARNING: PowerCLI is session based. If you exit the PowerCLI session without first exporting the bundle to the repository, the image cannot be reused.

- h. Add the deploy rule that you created.

```
add-deployrule -deployrule "auto-deploy-rule-name?"
```

8. Set up DHCP to network-boot the ESXi host:

- a. Get the MAC address of the ESXi host.

- b. On the DHCP server:

- Create an IP reservation for the ESXi host using its MAC address.
- Add option 66 (Boot Server Host Name - TFTP server IP).
- Add option 67 (Bootfile Name - undionly.kpxe.vmw-hardwired).

9. Boot the ESXi host.



NOTE: The ESXi host should appear in a vCenter data center automatically.

10. Verify that the vGW Series VIB was installed. Select the ESXi host, click the **Hardware Status** tab, and expand **Software Components** to ensure that dvfilter-altor-vf exists.

You must set up a host profile for the cluster where your hosts will be booted.

1. Before you create the host profile, set up the following components.

- The network and storage.
- Additional vNICs.

2. If you have multiple clusters, create a separate deploy rule for each ESXi host to direct the new host to a specific cluster. For example:

```
New-DeployRule -name "HostCluster" -item cluster-name -Pattern
"ipv4=10.70.1.1-10.70.1.250"
```

```
Add-DeployRule -DeployRule HostCluster
```

3. Use **Security configuration > Administrator password** to configure the administrator password after you create the host profile.

Configuring vGW Series Auto Deploy Support

Configure vGW Series to install a vGW Security VM on selected ESXi hosts that were deployed through VMware Auto Deploy. Use the Automatic Securing of Auto-deployed hosts pane on the Install Settings page. See [Figure 117 on page 232](#).

Figure 117: Configuring Automatic Installation of vGW Security VMs for Auto-Deployed ESXi Hosts

Automatic Securing of Auto-deployed Hosts

Select which auto-deployed hosts you would like to auto-secure

☐ No Host

☐ All Hosts

☒ Hosts in the following clusters:

☒ cluster_a

☒ cluster_b

☐ cluster_c2lk

The Security VM will be automatically named, using a prefix string you supply appended with the last octet of its corresponding ESX Host IP address, in the form of [prefix]_[octet].

SVM Name prefix: Example: SVM_123

Port Group:

Datastore:

Security VM IP address assignment: A static address will be created from the provided subnet followed by the last octet of the ESX IP address.

Method:

IP Address: . . . XXX

Network Mask:

Default Gateway:

☐ Force recheck on all hosts

Save

vGW Series automatically installs a vGW Series VM on the selected hosts.

1. Select the ESXi hosts to secure.
 - **No hosts**—No hosts will be secured.
 - **All hosts**—All hosts will be secured.
 - **Hosts in the following clusters**—Only hosts in the clusters that you identify will be secured. Select the check box for each cluster that you want to include.
2. Specify a prefix to use as part of the name that is assigned to automatically installed vGW Security VMs. Select the port group and the datastore to use.
 - **SVM Name prefix**—vGW Series automatically assigns a name to a vGW Security VM using the value you specify as the prefix. To create the complete name, it prepends the value to the last octet of the ESXi host IP address in the format [prefix]_[octet].

For example, if you used SVM_ as the prefix, if the last octet of the ESXi host IP address to secure was 123, the name SVM_123 would be assigned to the vGW Security VM for that host.

- **Port Group**—From the **Port Group** list, select the network label for the port groups.

Port groups serve as anchor points for VMs that connect to labeled networks. A port group is identified by a unique network label. The same network label is used for all port groups in a datacenter that are physically connected to the same network.

When you select either All Hosts or specific clusters, vGW Series updates the port group selection list. The list includes only port groups that are common to *all* connected hosts.

This behavior applies if you select one cluster or more than one.

- **Datastore**—From the Datastore list, select the datastore to use for the vGW Security VMs.

When you select either All Hosts or specific clusters whose hosts are to be secured, vGW Series updates the datastore list to include their datastores. The list includes only options that are common to *all* connected hosts. If there are no datastores that are on *all* hosts, the list is empty. However, if there is only one connected host, the list will show all of the datastores on that host.

This behavior applies to all clusters, whether you select one or more.

3. Select the method to use to acquire IP addresses for the vGW Security VMs.

- **Method**—Select either DHCP or static.
- **IP Address**—If Method is set to static, specify the static IP address to assign to the vGW Security VM.
- **Network Mask**—Specify the network mask to use in the IP address for the vGW Security VM.
- **Default Gateway**—Specify the default gateway for the vGW Security VMs.

4. Reset the error count to override the limit restricting the number of times that vGW Series is allowed to attempt to install a vGW Security VM on a host after repeated failures, reset the error count. Select **Force recheck on all hosts**.

vGW Series maintains a count of the number of failed attempts for each host. When that count is exceeded, it no longer tries to install a vGW Security VM on it. The installation attempts limit is set in the **center.auto.deploy.svm.install.retry.count** parameter, which has a default of three times. If you select this check box, the count is reset. It is also reset if you modify configuration settings.

You can create a per-host XML configuration. If you do this, the file must reside at `/usr/lib/tomcat/webapps/ROOT/WEB-INF/autoDeploy.xml`. You can find the xsd to use at:
<http://vgw-milford.juniper.net/trac/browser/center/branches/fullers/schemas/autoDeploy.xsd>.

If the static IP and IP or netmask or gateway are not set, the fallback behavior is to use configuration information set in the vGW Security Design VM.

- If the static IP and IP or netmask or gateway are not set, the fallback behavior is to use configuration information set in the vGW Security Design VM is used.

For example, if you set the IP method to DHCP in the vGW Security Design VM and a per-host configuration host entry does not have the IP configuration method specified, then the vGW Security VM for that host would get DHCP.

- If the port group or datastore are not found, the installation is canceled and a message is issued in the error log.

Understanding vGW Series Timeout Parameters and the vGW Security VM Installation, Uninstallation, and Update Tasks

Configurations for the following two timeout parameters affect a variety of vGW Security VM installation, uninstallation, and update processes:

- `center.timeout.vm.long.in.sec` (default: 10 minutes [600 seconds])
- `center.timeout.host.long.in.sec` (default: 10 minutes [600 seconds])

These JunosV Firefly Host VM processes entail individual tasks and groups of tasks. For example, the JunosV Firefly Host Module removal process that occurs when a JunosV Firefly Host VM is being uninstalled includes the "enter maint(enance) mode" and "remove fastpath" tasks.

If the ESX/ESXi host on which the vGW Security VM was installed exceeded the configured timeout value while it was being put into maintenance mode during the vGW Security VM uninstallation, the message that vGW Series reported prior to vGW Series 5.5 might have been misleading because it pertained to the *group* of tasks comprising the kernel module removal process.

Beginning with vGW Series 5.5, when a task exceeds the configured timeout value that pertains to it, vGW Series generates a log error entry that describes the individual task that was being executed when the timeout event occurred and the timeout parameter configuration that controls it, rather than giving a single task group message.

For example, the following message is generated and written to the log when the process of cloning the vGW Security VM template exceeds the amount of time configured for `center.timeout.vm.long.in.sec`.

"Cancelled task (cloning Security VM X from template Y) as it was taking too long.
Timeout set by center.timeout.vm.long.in.sec"

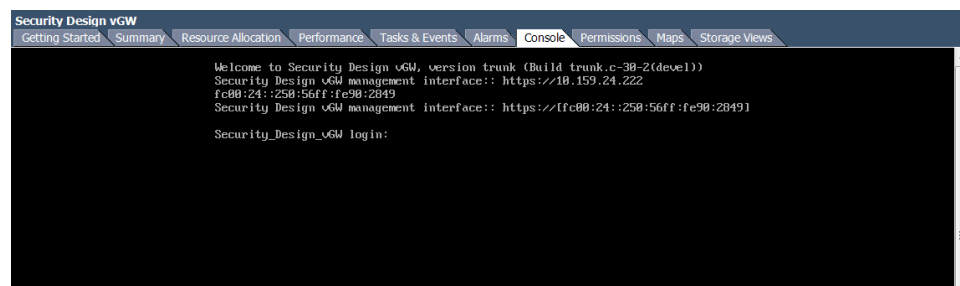
The timeout parameters are configurable to allow you to adapt your configuration to different vCenter behaviors. For example, a log entry might indicate that a vCenter task is taking longer than expected. You can use the console to run the vGW Series command-line interface (CLI) and change the configuration for the timeout parameter affecting the task. You can adjust the configuration appropriately and retry the process.

To use the vGW Series CLI from the vCenter console:

1. Launch the VMware vSphere Client.

2. Right-click the vGW Security Design VM icon on the left navigation panel to display a list of options.
3. Select the third option on the list, **Open Console**. Alternatively you can select the Console tab, as shown in [Figure 118 on page 235](#).

Figure 118: vGW Series CLI Console



The console window appears.

Some of the tasks affected by these timeout parameters are:

- vGW Security VM shutdown
- ESXi reboot
- cloning vGW Security VM template
- vGW Security VM reporting heartbeat with new version after update

Related Documentation

- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Removing vGW Security VMs from ESX/ESXi Hosts on page 235](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)
- [Updating the vGW Security Design VM on page 319](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding vGW Series on page 3](#)

Removing vGW Security VMs from ESX/ESXi Hosts

This topic explains how to remove a vGW Security VM from an ESX or an ESXi host.

If you want to remove the VMX entries before you un-install the vGW Security VM, then before unsecuring the entire host by removing the vGW Security VM, unsecure the individual VMs. See [“Securing and Unsecuring Virtual Machines Using the vGW Security Design VM” on page 237](#).

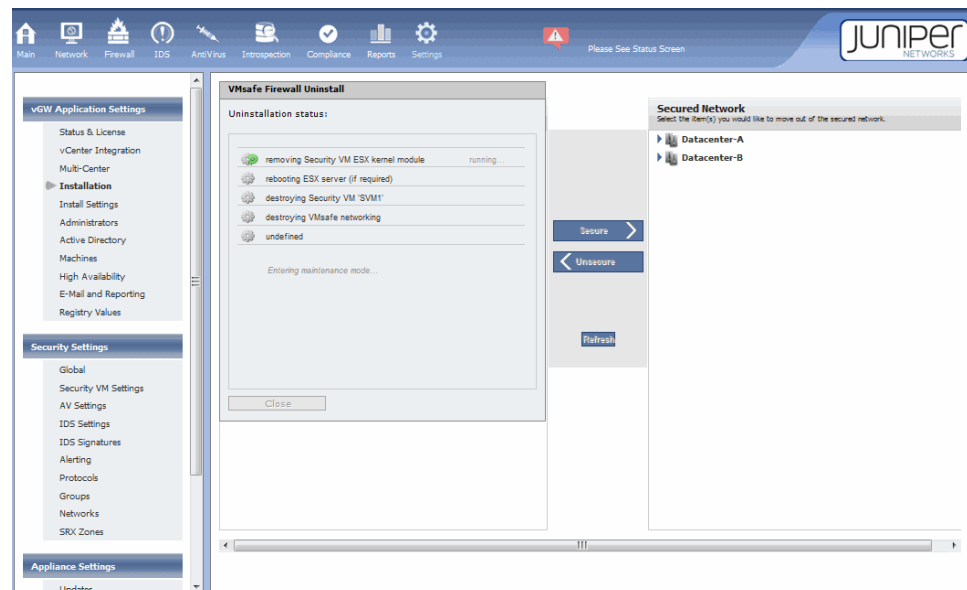
To un-install the vGW Security VM from a host:

1. In the Secured Network pane of the Settings module vGW Application Settings > Installation page, select the host that you want to move out of the secured network.

2. Click the **Unsecure** arrow button.
3. The VMsafe Firewall Uninstall status pane is displayed. As the vGW Security Design VM removes the firewall from the host—or moves a specific VM out of the secured network, if you selected a VM—the status pane identifies the active process.

When you select an individual VM to remove from the secured network and click **Unsecure**, the vGW Security Design VM removes all relevant VMX entries for that VM, reverting the VM to its state prior to vGW Series protection of it. [Figure 119 on page 236](#)

Figure 119: vGW Security VM Uninstall



If you plan to un-install vGW Series from your virtualized environment, unsecure all VMs in this manner. Afterward, select the check box for each of the ESX/ESXi hosts and click **Unsecure** to remove them from vGW Series protection. This process removes the kernel module and the related VMservice vSwitch and port groups.

Unsecuring a host before removing its VMs does not affect the VMs adversely. However, the process does not remove VMsafe VMX entries that pertain to vGW Series. These entries are no longer required by that VM.



NOTE: You might not want the VMX entries for a VM to be removed under these conditions. For example, you might want to remove only the vGW Series kernel module from a specific host. This might be the case if you want the VMs to be moved to a different ESX/ESXi host for protection, or you intend to reinstall vGW Series later.

Related Documentation

- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Understanding vGW Series on page 3](#)

Securing and Unsecuring Virtual Machines Using the vGW Security Design VM

After you install the vGW Security VM on an ESX/ESXi host to secure it, the vGW Security Design VM allows you to manually secure virtual machines (VM) on that host or remove them from the protected network. Removing a secured VM from the protected network is referred to as *unsecuring* the VM.

To secure a VM that does not belong to the Secured Network:

1. In the vGW Security Design VM Settings module vGW Application Settings section, select **Installation**.
2. In the Unsecured Network pane, select the VM that you want to secure. Click the check box in front of its name.
3. Click **Secure**.

As it secures the VM, the vGW Series reports on the status of each part of the process. If the VM is successfully secured, the report states that the VM was successfully secured.

4. Click **Close**.

The vGW Security Design VM displays a process symbol that dynamically indicates that the VM is being secured with a firewall and moved into the secured network. The VM is now protected, and it appears in the Secured Network pane.

After all vGW Series components in your environment are upgraded to release 5.5, if you attempt to introduce components from a previous release, the process is halted and vGW Series displays a message informing you that you must install the correct version.

Related Documentation

- [Disabling the vGW Series Suspend-Resume Process Enacted After a VM Is Unsecured on page 237](#)
- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Understanding vGW Series on page 3](#)

Disabling the vGW Series Suspend-Resume Process Enacted After a VM Is Unsecured

You use the vGW Security Design VM Settings module Installation section to secure and unsecure a VM. By default, the vGW Series suspends and resumes a VM when you unsecure it. You can change this behavior by changing the value of the `vm-safe.config` option.

- [Displaying the State of the `vm-safe.config` Setting on page 237](#)
- [Disabling the Suspend-Resume Process on page 238](#)

Displaying the State of the `vm-safe.config` Setting

This example shows the default setting. You can use the following command to display the current state of the `center.config vm-safe.config` option:

```
(Cmd) config show center.suspend.after.vmsafe.config
# whether center should suspend and resume VM after VMSafe configuration
center.suspend.after.vmsafe.config = true
```

Disabling the Suspend-Resume Process

In some cases it might be necessary or desirable to stop vGW Series from enacting the suspend-resume process after a VM is unsecured. For example, you might want to disable the process to allow the VM to be migrated to another host or to suspend and resume the VM later after completing the removal of protection from the VM.



TIP: Take care when you protect VMs such as the VMware vCenter Database VM and other VMs that must not be suspended.

To enable the vmsafe config process to take effect after the VM is migrated to another host without suspending the VM, use the following statement. Set the option to false in center.config:

```
(Cmd) config set center.suspend.after.vmsafe.config false
```

After changing this value, either restart the vGW management process or reboot the vGW Security Design VM. You can use the service restart command line or the vGW Security Design VM to restart the vGW management process.

To restart the vGW management process from the command line, enter the following command:

```
(Cmd) service restart tomcat
Sending 'restart' command
The following watches were affected:
tomcat
```

To restart the vGW management process using the vGW Security Design VM:

1. Select the Settings module Support section.
2. In the Restart pane of the displayed page, click **Restart**.

Related Documentation

- [Installing vGW Security VMs on ESX/ESXi Hosts on page 218](#)
- [Securing and Unsecuring Virtual Machines Using the vGW Security Design VM on page 237](#)
- [Understanding vGW Series on page 3](#)

Understanding Automatic Securing of VMs

vGW Series allows you to configure your system to *automatically* secure VMs. Auto-securing VMs streamlines policy application allowing you to efficiently ensure security throughout your virtual infrastructure. You can configure the Auto-Secure feature options to direct vGW Series to automatically secure VMs in the manner most appropriate for your environment.

You use the Settings module vGW Application Settings > Install Settings > Automatic Securing of VMs pane to configure Auto-Secure for your virtualized environment.

The Automatic Securing of VMs pane includes the following options:

- No VM

No individual VMs or groups of VMs are automatically secured. This is the default behavior.

- VMs in the following group

This option allows you to select either a Static Group or a Smart Group from the list of existing groups. The list contains all groups, including those configured as Policy Groups and those that are not. Using this option, you can select only one group.



NOTE: Only VMs in the selected group are automatically secured.

- If you did not configure the selected group as a Policy Group, vGW Series automatically secures members of the group with the Global and Default policies.
- If you configured the selected group with the Policy Group option, then any policy rules that were created for the group and applied to it take effect. In this case, the Default policy is not used.

- VMs with a VM Policy or in a Policy Group

Because Default Policy and Global Policy rules tend to be restrictive, they are not appropriate for securing all VMs. This option allows you to predefine policy rules for individual VMs and groups of VMs and direct vGW Series to use the policy rules that you predefined to automatically secure them rather than relying on just the Default and Global policy rules. Using this option, you can automatically secure many Policy Groups and individual VMs instead of being restricted to selecting a single group.

VMs that fit any of the following criteria are automatically secured:

- Individual VMs for which you have predefined specific policy rules and applied those policies using the Firewall module Apply Policy page to install the policy.
- Groups of VMs that you created as Static Groups or Smart Groups and for which you selected the Policy Group option. You must also have created and applied a policy for the group, and that policy must contain rules.

- All VMs

All VMs are automatically secured. As described previously, any policy rules defined for Policy Groups that have been previously applied take effect for VM members of the group. If a VM is not a member of any group, then Global and Default Policies and any individual VM rules take effect for them.

You can refine this selection by excluding a specific group of VMs.

- Optionally, exclude a group of VMs from being automatically secured. You might want to exclude VMs from auto-securing that you are using for testing.



NOTE: vGW Series auto-secure feature will not attempt to secure an FT-enabled VM. vGW generates an alert telling you that you must disable FT for that VM or suspend the VM for vGW to secure the VM. The auto-secure feature monitors for cases in which an FT-enabled VM is disabled and for VMs that are suspended and powered-off.

If a VM is automatically secured, you cannot use the Settings module Installation page to unsecure it. The VM is shown on this page in a dimmed box and a message is presented informing you that it is automatically secured. In this case, if you were able to unsecure the VM, vGW Series would simply secure it again automatically.

Instead, you must first remove the VM from the automatically secured group that it belongs to, or, if it is an individual VM, remove the policy from it, and then unsecure it.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)

Understanding the vGW Series Split-Center Feature

This topic covers the vGW Series Split-Center feature that allows you to segment resources contained in a single VMware vCenter into multiple domains, or scopes, independently managed by different vGW Security Design VMs. The Split-Center feature allows for improved resource isolation for cloud services and multi-tenancy. It supports unlimited scalability as the VMware vCenter capacity increases and your deployment takes advantage of it. You can deploy as many vGW Security Design VMs as are needed as you scale your environment.

Together the individual vGW Security Design VMs associated with a vCenter can collectively secure all its ESX/ESXi hosts and VMs, but each individual vGW Security Design VM manages only a specific set of resources, determined according to how you configure the management scope for that vGW Security Design VM. One vGW Security Design VM does not have visibility into another vGW Security Design VM or the parts of the virtualized environment that another vGW Security Design VM secures. After you configure its management domain, to a single vGW Security Design VM it is as if all objects outside its scope do not exist.

The Split-Center feature allows you to configure management domains that consist of:

- entire vCenter

You can select the entire vCenter as the management scope. Effectively you are not using the Split-Center feature in this case.

- Multiple data centers

Ensure that each data center is assigned to only one vGW Security Design VM. Otherwise, unexpected consequences can occur.


- One or more clusters of hosts within a data center

In some environments, organizations use clusters of host to segment their virtual infrastructure rather than data centers. To support these environments, you can configure the Split-Center feature along lines of host cluster management domains.



NOTE: All of the host clusters that you select to belong to a management domain (scope) must be in the *same* data center. You can not include host clusters from two or more different data centers in the scope.

The following figure shows a vCenter with two data centers. Nine of the hosts in Data center 1 comprise a cluster that is configured as a domain to be secured and managed by vGW Security Design VM1. The remaining hosts in data center 1 are not secured by vGW Series. All of the hosts in Data center 2 are configured as a single domain to be managed and secured by vGW Security Design VM2.



As the administrator who oversees your deployment, most likely you determine the management domains for your virtualized environment. Afterward, you can convey to administrators of the individual vGW Security Design VMs for the vCenter which objects to include in their management scopes.

Administrators of various vGW Security Design VMs establish their management domains when they run the vGW Installation Wizard to initially set up their vGW Security Design VM. See [“Setting Up vGW Series” on page 36](#) for details on initially setting the management domain.

The following figure shows how to configure the management scope, or domain during installation.

Figure 120: Configuring the Management Scope During Installation to Include Clusters

If you change the management domain deployment design later, administrators can reconfigure the Split-Center domains that their vGW Security Design VMs manage using the Settings module vGW Application Settings > vCenter Integration > vCenter Settings pane called *Select a scope for your Security Design vGW*.

For both the initial management domain configuration during product installation and when you change the configuration, the same vCenter Settings pane is used, but it is approached differently.

Depending on how you define your management domain, you configure a vGW Security Design VM for Split-Center in either of the following ways:

- **Datacenter**—A subset of data centers in the vCenter.

In this case, the vGW Security Design VM is able to access and manage only the VMs and other entities in the selected data centers.

To use this scope:

1. Select **Datacenters**.

vGW Series displays all of the vCenter's data centers.



NOTE: To update the list of data centers at any time to show changes—datacenters that might have been added or removed—click **Refresh**.

2. Select the data centers for your vGW Security Design VM to manage.

Ensure that each data center is assigned to only one vGW Security Design VM. Otherwise, unexpected consequences can occur.

3. Click **Save**.

- **Clusters**—A subset of host clusters in a data center. In this case, the vGW Security Design VM is able to access only the VMs and other entities on the selected host clusters.

To use the Clusters scope:

1. Click **Clusters** in the *Select a scope for your Security Design vGW* area.
vGW Series displays a list of available data centers in response.
2. Select a data center from the displayed list whose host clusters you want the vGW Security Design VM to manage.
 - a. Click the arrow at the end of the box beside **Datacenter:** to display a list of data centers for the vCenter.
 - b. Click the data center whose cluster(s)/host(s) you want your vGW Security Design VM to manage.
vGW Series displays a list of cluster(s)/host(s) for the data center that you selected.
3. Select the check box before the names of the cluster(s)/host(s) that you want to include in your management domain. All host clusters that you select must belong to the same data center.
4. Click **Save**.



NOTE: Ensure that each host cluster is assigned to only one vGW Security Design VM. Otherwise, unexpected consequences can occur.

You can change the cluster selection at any time. However, when you change the cluster scope, either of the following conditions can occur:

- Some vGW Security VMs could become unmanaged—This can occur when you remove a cluster from the list of selected clusters. Any vGW Security VM installed on an ESX/ESXi host that belongs to the removed cluster will no longer be accessible, and therefore it is no longer managed by the vGW Security VM.
- Some unmanaged vGW Security VMs could become accessible—If ESX/ESXi hosts that belong to a cluster that you add to your vGW Security Design VM management domain had a vGW Security VM installed on them by a different vGW Security Design VM, you could gain access to the vGW Security VMs. It is possible and important to gain access to an unmanaged vGW Security VM when you add its host cluster to your vGW Security VMs management domain for the following reason.

When a vGW Security VM becomes inaccessible because of cluster or datacenter selection changes its original vGW Security Design VM, its operational state might be compromised unless it is imported into another vGW Security Design vGW. This is because the vGW Security VM continues to try to communicate with its original vGW Security Design VM, which no longer recognizes it as a managed.

To view a list of unmanaged SVMs and render them manageable again:

1. Display the Settings module > Security VM Settings page.

The unmanaged vGW Security VMs are identified by a gray triangle status indicator.

2. To make a vGW Security VM manageable again, click its row to select it.
3. Click **Import**.

- Clusters

To configure a domain that contains clusters:

1. Select **Clusters** in the *Select a scope for your Security Design vGW* area.
2. Select a data center.
 - a. Click the arrow at the end of the box beside **Datacenter:** to display a list of data centers for the vCenter.
 - b. Click the data center whose cluster(s)/host(s) you want your vGW Security Design VM to manage.

vGW Series displays a list of cluster(s)/host(s) for the data center that you selected.
3. Check the boxes before the names of the cluster(s)/host(s) that you want to include in your management domain.

You can define the management scope initially using the vGW Installation Wizard when you set up your vGW Security VM. You can use the same page later to change the configuration by selecting the Settings module vGW Application Settings > vCenter Integration > vCenter Settings page *Select a scope for your Security Design vGW* area.



NOTE: When you configure the management domain scope, you can select the entire vCenter. Effectively, Split-Center is not used in this scenario because a single vGW Security Design VM is responsible for all data centers in the vCenter.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Configuring Scaling Using the Multi-Center and Split-Center Features on page 251](#)
- [Configuring vGW Series Multi-Center on page 247](#)
- [Understanding the vGW Security Design VM on page 67](#)

Understanding the Multi-Center Feature

This topic covers the vGW Series Multi-Center feature that synchronizes policy across vGW Security Design VM management centers to enable large scale virtualization. The

Multi-Center feature is useful for large-scale virtualized environment deployments spread across many vCenters.

This section includes the following sections:

- [The Multi-Center Feature on page 245](#)
- [Deploying vGW Series in an Environment With a Mix of Delegate and Stand-alone vGW Security Design VMs in Various vCenters on page 246](#)

The Multi-Center Feature

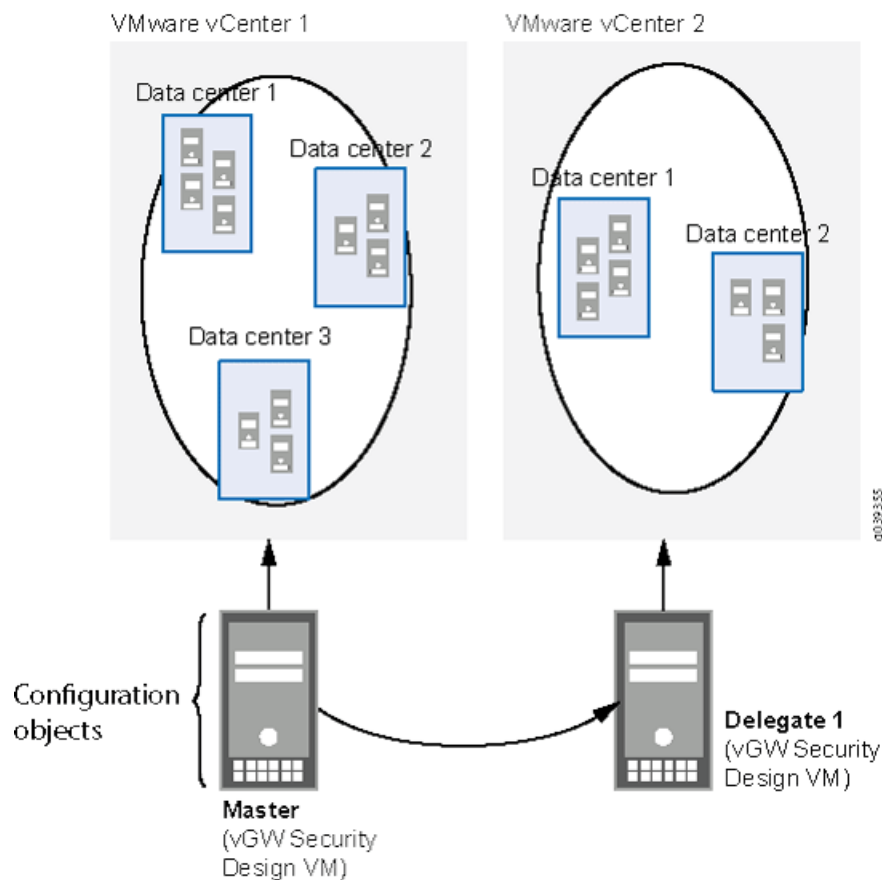
For various reasons—such as geographic separation of data centers, scaling requirements, and use of different administrative domains—some companies who deploy the vGW Series must use more than one VMware vCenter to manage their environments. These companies want to use the same or similar vGW Security Design VM configuration for all of their data centers, as if they were rolling out a single deployment. Manually configuring separate vGW Security Design VMs at various locations with the same information consumes time, and it is cumbersome and error prone.

To accommodate companies with these requirements and companies that want to scale their environments for other reasons, the vGW Series includes a feature called Multi-Center. The Multi-Center feature allows you to designate a single vGW Security Design VM connected to a vCenter at one location as the master.

Following the database replication model, configuration is done at master vGW Security Design VM. It can be synchronized all or in part to one or more delegate vGW Security Design VM centers, each of which is connected to an individual vCenter. Configuration of global objects at the master vGW Security Design VM is propagated to the delegate vGW Security Design VMs centers automatically, based on objects selected when the administrator of the master vGW Security Design VM creates a Multi-Center definition for the delegate center.

[Figure 121 on page 246](#) shows a master center and a delegate center. Objects at the master center are synchronized to the delegate center.

Figure 121: vGW Series Multi-Center



Deploying vGW Series in an Environment With a Mix of Delegate and Stand-alone vGW Security Design VMs in Various vCenters

The vGW Security Design VM Multi-Center feature can be in whatever configuration your environment requires. You might design your virtualized environment to include some vGW Security Design VMs that belong to a configuration that uses the Multi-Center feature and some that do not. You might want one vGW Security Design VM to manage resources at a specific vCenter and let it have an entirely unique configuration. You might want others at different vCenters to use largely the same configuration.

For example, an organization's virtualized environment might include six data centers of various sizes, each of which is connected to an individual vCenter. The administrator uses the same overall configuration for five of the data centers but not for the sixth one. The Multi-Center feature suits this environment well also in that it can secure the five data centers in the same way, but the administrator of the vCenter environment with different security requirements could define his own policies and other security protection independently.

Related Documentation

- [Configuring vGW Series Multi-Center on page 247](#)

Configuring vGW Series Multi-Center

This topic explains how to configure the Multi-Center feature which allows you to synchronize the configuration at one vGW Security Design VM across multiple vGW Security Design VMs connected to different VMware vCenters. The Multi-Center feature allows you to streamline configuration across multiple vGW Security Design VMs and coordinate various aspects of security as you scale. It relies on the configuration at one vGW Security Design VM, referred to as the master center, which is synchronized in part or whole to other vGW Security VMs, referred to as delegate centers.



NOTE: You can also use Multi-Center with the Split-Center feature to synchronize the configuration across multiple vGW Security Design VMs that manage resources in the same vCenter.

Before you read this topic, read [“Understanding the Multi-Center Feature” on page 244](#).

This topic contains the following sections:

- [vGW Security Design VM Master Center on page 247](#)
- [vGW Security Design VM Delegate Centers on page 247](#)
- [Configuring Multi-Center on page 248](#)

vGW Security Design VM Master Center

As administrator of the master center, you configure the object synchronization for all delegate centers—at the master vGW Security Design VM. After you configure the vGW Security Design VM that you will use as the master center, you can define delegate center configurations for individual delegate centers.

Although you configure Multi-Center for all delegate centers, each delegate center has its own independent configuration, and they can differ. When you add a delegate center configuration, you designate the objects that are synchronized to it.



NOTE: The master vGW Security Design VM and the delegates must be able to communicate using addresses from the same IP protocol family. Communication problems should not exist if this is the case. Too, if either of them is configured for dual stack, problems should not exist. If both are configured with a single IP from different protocol families, problems could ensue. To solve this problem, you could change the IP address used for one of them.

vGW Security Design VM Delegate Centers

Administrators of the master and delegate centers cooperate in implementing Multi-Center. They determine the objects to synchronize to the delegate center from the master center. Each delegate center has its own configuration at the master center.

Configuration objects, such as policies, configured at the master center that are synchronized to delegate centers are viewed as global objects from the perspective of the delegate center. Some delegate centers might not synchronize a certain object, but rather retain their own local configuration for that object. For information about configuration objects and how they are synchronized, see [“Understanding vGW Series Multi-Center Synchronized Objects” on page 250](#).

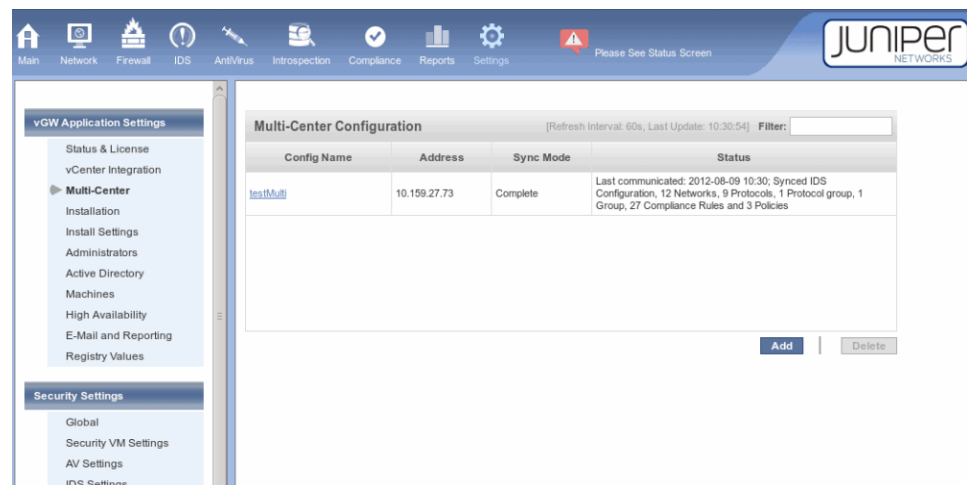
A vGW Security Design VM delegate center is created for a vCenter no differently from how it would be if it were independent. You import the OVA into the vCenter to be secured. For information on how to integrate vGW Series with vCenter, see [“Using the OVA Bundled Method to Integrate vGW Series with the VMware Infrastructure” on page 24](#). After the installation is complete, you can begin to engage the vGW Security VM in the Multi-Center configuration.

Configuring Multi-Center

To configure Multi-Center, use the Settings module vGW Application Settings > Multi-Center page on the master center. To add a delegate center to the Multi-Center configuration:

1. At the bottom Multi-Center Configuration pane, click **Add**. See [Figure 122 on page 248](#).

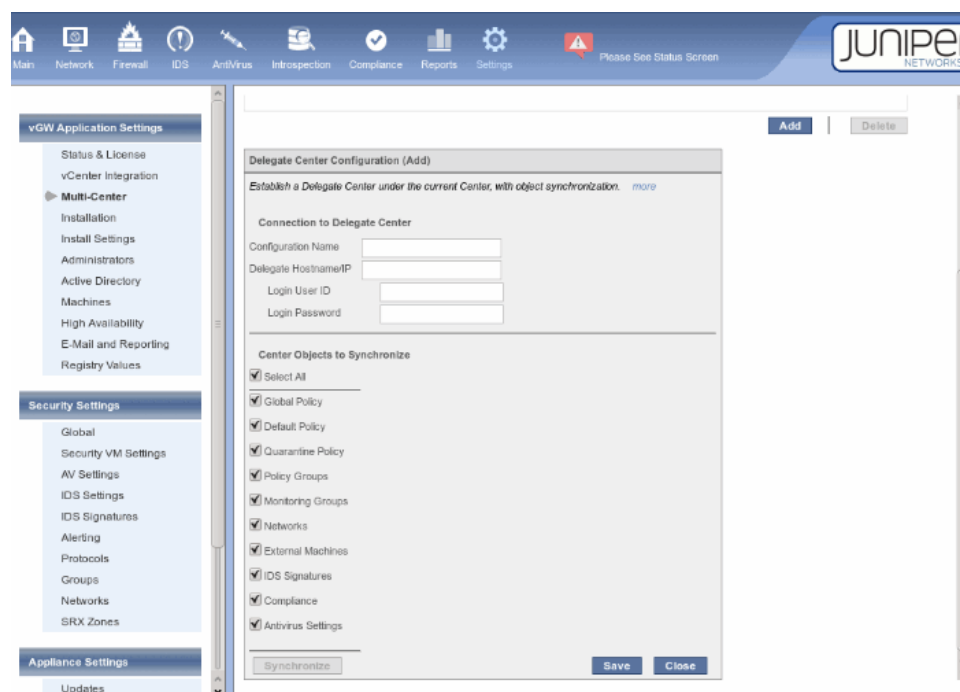
Figure 122: Multi-Center Configuration Page at Master vGW Security Design VM



The Delegate Center Configuration (Add) pane is displayed on the master vGW Security Design VM.

2. In the **Configuration Name** field, specify a name for the configuration that represents the delegate center. Note that the name field is used only for reference, and it can be anything. It does not need to match the name of the delegate vGW Security Design VM. See [Figure 123 on page 249](#).

Figure 123: Delegate Center Configuration on the Master vGW Security Design VM



3. In the **Delegate Hostname/IP** field, enter the name or the IP address of the delegate center.

Enter a valid hostname, IPv4 address, or IPv6 address.
4. In the **Login User ID** and **Login Password** fields, enter the delegate center's authentication information.
5. In the **Center Objects to Synchronize** pane, select the objects to synchronize.
 - Check **Select All** if you want the state of all of the objects in the list to be synchronized from the master vGW Security Design VM to the delegate center that you are defining.
 - If you want only some of the objects to be synchronized from the master vGW Security Design VM to the delegate center, select the check box before each object to synchronize.
 - Global Policy—Synchronizes the global policy and all objects it depends on. Among other objects, configurations for the source and destination of the rules in the policy and the protocols are copied.
 - Default Policy—Synchronizes the default policy and all objects it depends on. Among other objects, configurations for the source and destination of the rules in the policy and the protocols are copied.
 - Quarantine Policy—Synchronizes the quarantine policy and all objects it depends on. Among other objects, configurations for the source and destination of the rules in the policy and the protocols are copied.

- Policy Groups—Synchronizes all the policy groups and policies associated with them, and all objects that they depend on. Among other configurations, the sources and destinations of the rules in the policies, the protocols, the networks and the machines in the groups are copied.
- Monitoring Groups—Synchronizes all the monitoring groups and the policies associated with them, and all objects that they depend on. Among other configurations, the sources and destinations of the rules in the policies, the protocols, the networks and the machines in the groups are copied.
- Networks—Synchronizes all networks.
- External Machines—Synchronizes all external machines.
- IDS Signatures—Synchronizes IDS Signatures and Settings.
- Compliance - Synchronizes compliance rules and all objects that they depend on, such as groups.
- Antivirus Settings—Synchronizes all AntiVirus scan configurations, and all objects that they depend on, such as groups.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series Multi-Center Synchronized Objects on page 250](#)

Understanding vGW Series Multi-Center Synchronized Objects

This topic explains how protocols and compliance rules objects are synchronized from the master vGW Security Design VM to the delegate vGW Security Design VM center. The Settings pane of the delegate center shows status and other information about objects that are synchronized to it.

From the perspective of a delegate center, the synchronized objects are viewed as read-only global objects, and they cannot be modified.

This topic includes the following sections:

- [Object Synchronization on page 250](#)
- [Object Naming on page 251](#)
- [Creation of Objects Local to the Delegate vGW Security VM on page 251](#)

Object Synchronization

It can occur that a newly synchronized global object is identical in name and content to a local object on the delegate center. In this case, for global objects that contain default values such as protocols and compliance rules, the local object is converted to a global one. The global version of the local object on the delegate center vGW Security Design VM is marked as converted. All references to the local object are preserved, but now they pertain to the global object. Because the converted object is now a global object, it is accessible as a read-only object on the delegate center vGW Security Design VM. That is, the administrator of the delegate center vGW Security Design VM cannot modify it.

When an object is no longer mirrored, it is deleted from the delegate center unless it is used by local objects. That is, if it was converted from a local object such as a protocol, it is converted back to the local object at that time.

Object Naming

To avoid naming issues and Smart Group logic problems, when the same name for a global object and a local object exists in the same context, the global object takes precedence and the name is used for it. vGW Series marks the object as global, as viewed from the delegate center. The object with the conflicting name is renamed with the word local appended to it.

The administrator of the master vGW Security Design VM can remove an object from selection for a delegate center. In this case, the object is no longer a global one on the delegate center. If a local counterpart exists, it is now reinstated and the delegate center administrator can edit it.

Creation of Objects Local to the Delegate vGW Security VM

Administrators of delegate vGW Security Design VMs centers are still able to configure local objects for their own systems. These local objects remain local, and they have no affect on the master vGW Security Design VM configuration, with some exceptions. For example, the priority of local policy groups is always lower than global ones.

Related Documentation

- [Understanding the Multi-Center Feature on page 244](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)

Configuring Scaling Using the Multi-Center and Split-Center Features

This topic explains how to use the vGW Series Split-Center and Multi-Center features together to secure your virtualized environment as you scale.

These features are typically used together to:

- Allow for partitioned management of resources among multiple vGW Security Design VMs at an individual vCenter.

The Split-Center feature allows you to segment responsibility for portions of your resources at an individual vCenter among multiple vGW Security Design VMs. It is as if each vGW Security Design VM were connected to an individual vCenter.

For background on the Split-Center feature, read [“Understanding the vGW Series Split-Center Feature” on page 240](#).



CAUTION: When you configure the Split-Center feature, ensure that each data center is assigned to only one vGW Security Design VM. Otherwise, unexpected consequences can occur.

- Deploy largely the same configuration to all vGW Security Design VM delegate centers, including those that share responsibility for a single vCenter.

The Multi-Center feature facilitates configuration management as you scale your environment. You can use it to create configurations that are largely the same for vGW Security Design VMs at different vCenters and for vGW Security Design VMs sharing security management responsibility for resources at the same vCenter. You can effectively deploy the same configuration to them automatically with real-time updates.

For background on the Multi-Center feature, see [“Understanding the Multi-Center Feature” on page 244](#).

This topic contains the following sections:

- [vGW Series Split-Center Multi-Center Configuration Requirements on page 252](#)
- [About the Example on page 253](#)
- [Configuring Split-Center and Multi-Center for vGW Security Design VMs on page 255](#)

vGW Series Split-Center Multi-Center Configuration Requirements

This example addresses a customer environment with a virtualized infrastructure that includes data centers at three individual VMware vCenters:

- The first vCenter, vCenter1, includes five customer data centers. vCenter1 is located in Dallas, Texas. One data center is considerably larger than the others.

The customer uses the Split-Center feature to partition management of the vCenter1 data centers among two vGW Security Design VMs in the following way:

- vGW Security Design VM-1 manages the large data center, vCenter1-data-center-1.
- vGW Security Design VM-2 manages the other four data centers:
 - vCenter1-data-center-2.
 - vCenter1-data-center-3.
 - vCenter1-data-center-4.
 - vCenter1-data-center-5.
- The second vCenter, vCenter2, includes two customer data centers. vCenter2 is located in Minneapolis, Minnesota. vGW Security Design VM-3 manages both:
 - vCenter2-data-center-1.
 - vCenter2-data-center-2.
- The third vCenter, vCenter3, includes two data centers. vCenter3 is located in Raleigh, North Carolina. vGW Security Design VM-4 manages both:
 - vCenter3-data-center-1.
 - vCenter3-data-center-2.

About the Example

This customer's virtualized environment spans three vCenters at various locations. The customer plans to use the Split-Center feature to divide security management responsibility for resources at one of the vCenters among two vGW Security Design VMs.

The customer plans to deploy largely the same configuration for all vGW Security Design VMs. Because manually creating separate configurations with the same parameters is time consuming and error prone, the customer decides to use the Multi-Center feature to solve this problem.

The Multi-Center feature allows the customer to use a single vGW Security Design VM as the master center. Its configuration is copied to all slave, or delegate, vGW Security Design VMs.

For this example, vGW Security Design VM-3 serves as the primary center. The administrator of vGW Security Design VM-3 configures the Multi-Center feature for all delegate centers.

Using the Settings module Application Settings > Multi-Center, the administrator defines an entry for each delegate vGW Security Design VM center. For this example, delegate centers include:

- vGW Security Design VM-1

The configuration specifies that all objects are to be copied.

- vGW Security Design VM-2

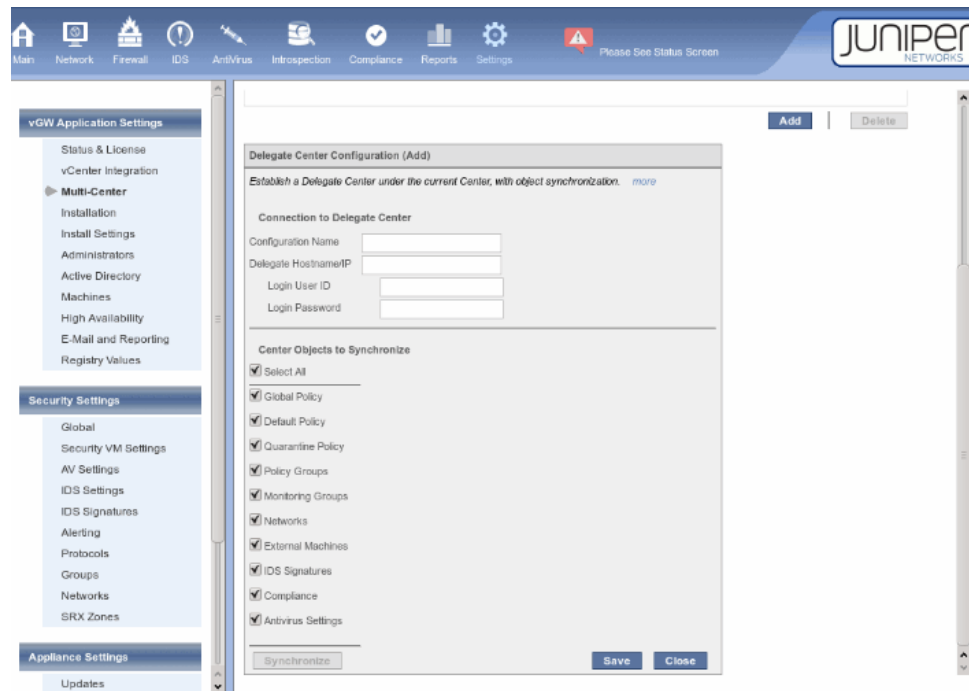
The configuration specifies that all objects are to be copied.

- vGW Security Design VM-4

The configuration specifies that all objects excluding monitoring groups and IDS are to be copied.

You use the Delegate Center Configuration (Add) pane of the Settings module Multi-Center feature to create an entry for a delegate vGW Security Design VM center. See [Figure 124 on page 254](#).

Figure 124: Delegate Center Configuration on the Master vGW Security Design VM



To do so, you provide the following information:

- In the **Configuration Name** field, specify a name for the configuration that represents the delegate center.



NOTE: Note that the name field is used only for reference, and it can be anything. It does not need to match the name of the delegate vGW Security Design VM.

- In the **Delegate Hostname/IP** field, enter the name or the IP address of the delegate center. This allows the master vGW Security Design VM and the delegate center vGW Security Design VM to communicate.
- In the **Login User ID** and **Login Password** fields, enter the delegate center's authentication information.
- In the **Center Objects to Synchronize** pane, select the objects to synchronize.
 - Check **Select All** if you want the state of all of the objects in the list to be synchronized from the master vGW Security Design VM to the delegate center that you are defining.
 - **Global Policy**—Synchronizes the global policy and all objects it depends on. Among other objects, configurations for the source and destination of the rules in the policy and the protocols are copied.

- **Default Policy**—Synchronizes the default policy and all objects it depends on. Among other objects, configurations for the source and destination of the rules in the policy and the protocols are copied.
- **Quarantine Policy**—Synchronizes the quarantine policy and all objects it depends on. Among other objects, configurations for the source and destination of the rules in the policy and the protocols are copied.
- **Policy Groups**—Synchronizes all the policy groups and policies associated with them, and all objects that they depend on. Among other configurations, the sources and destinations of the rules in the policies, the protocols, the networks and the machines in the groups are copied.
- **Monitoring Groups**—Synchronizes all the monitoring groups and the policies associated with them, and all objects that they depend on. Among other configurations, the sources and destinations of the rules in the policies, the protocols, the networks and the machines in the groups are copied.
- **Networks**—Synchronizes all networks.
- **External Machines**—Synchronizes all external machines.
- **IDS Signatures**—Synchronizes IDS Signatures and Settings.
- **Compliance**—Synchronizes compliance rules and all objects that they depend on, such as groups.
- **Antivirus Settings**—Synchronizes all AntiVirus scan configurations, and all objects that they depend, such as groups.

Configuring Split-Center and Multi-Center for vGW Security Design VMs

Configuring Split-Center for the First vGW Security Design VM

Step-by-Step Procedure This configuration shows how to use the Split-Center feature to give vGW Security Design VM-1 management responsibility for part of the resources at vCenter1.

From the Settings module **vGW Application Settings** > **vCenter Integration** page:

1. In the vCenter Settings pane, enter the following information:
 - The server name or IP address of the vCenter. For this example, enter **vCenter1**.
 - The vGW Security Design VM-1 username and password to authenticate to vCenter1. For this example, enter **admin-1** and **talk#321**.
2. In the vCenter Settings pane, select a management scope for vGW Security Design VM-1. To display the data centers belonging to vCenter1, select the **Selected Datacenters** option button.

The data centers belonging to vCenter1 are displayed:

- vCenter1-data-center-1
- vCenter1-data-center-2
- vCenter1-data-center-3

- vCenter1-data-center-4
- vCenter1-data-center-5

By default, the system is configured to allow the vGW Security Design VM to manage all data centers.

3. Click the check box before vCenter1-data-center-1, and click **Save** to allow vGW Security Design VM-1 to manage it.

vGW Security Design VM-1 will now be able to manage only the VMs and other resources for vCenter1-data-center-1 of vCenter1.



NOTE: Before the system saves your selection, vCenter1 verifies the authentication credentials that you specified. The system displays the following message:

Checking vCenter login credentials. This may take up to 15 seconds depending on server loads.

If your credentials are invalid, your data center scope management selection is not committed.

4. If you want to commit the configuration, click **Okay**.

Configuring Split-Center for the Second vGW Security Design VM

Step-by-Step Procedure

This configuration shows how to use the Split-Center feature to give vGW Security Design VM-2 management responsibility for part of the resources at vCenter1.

1. From vGW Security Design VM-2, select the Settings module.
2. In the navigation tree, select vCenter Integration beneath vGW Application Settings.
3. In the vCenter Settings pane, enter the following information:
 - The server name or IP address of the vCenter. For this example, enter **vCenter1**.
 - The vGW Security Design VM-2 username and password to authenticate to vCenter1. For this example, enter **admin-2** and **talk#4*5#6**.
4. In the vCenter Settings pane, select a management scope for vGW Security Design VM-2. To display the data centers belonging to vCenter1, select the **Selected Data centers** option button.

The data centers belonging to vCenter1 are displayed:

- vCenter1-data-center-1
- vCenter1-data-center-2
- vCenter1-data-center-3

- vCenter1-data-center-4
- vCenter1-data-center-5

By default, the system is configured to allow the vGW Security Design VM to manage all data centers.

5. Click the check boxes before vCenter1-data-center-2, vCenter1-data-center-3, vCenter1-data-center-4, vCenter1-data-center-5, and click **Save** to allow vGW Security Design VM-2 to manage them.



NOTE: Before the system saves your selection, vCenter1 verifies the authentication credentials that you specified. The system displays the following message:

Checking vCenter login credentials. This may take up to 15 seconds depending on server loads.

If your credentials are invalid, your data center scope management selection is not committed.

6. To commit the configuration, click **Okay**.

Defining Entries for a Delegate Center Using the Multi-Center Feature

Step-by-Step Procedure

This example shows how to define entries for one of the three vGW Security Design VMs to allow it to become a delegate center and inherit most of the vGW Security Design VM-3 master's configuration. Configuration of the other two delegate centers is not shown here, but it is done similarly to the single configuration example.

This example shows how to configure:

- Entries for vGW Security Design VM-1 and vGW Security Design VM-2 to allow all configuration objects to be copied to them.
- An entry for vGW Security Design VM-4 to allow all configuration objects excluding monitoring groups and IDS to be copied to it.

To define a delegate center entry for vGW Security Design VM-1, from the vGW Security Design VM-3 master Settings module **vGW Application Settings > Multi-Center** page:

1. Enter **mc-delegate-1** as the name for the delegate center entry.
2. Enter **admin-1** and **talk#321** as the user ID and password credentials of the delegate center.
3. Under Synchronize Objects, click **Select All**.
4. If you are satisfied with the configuration, click **Save**. Otherwise, click **Cancel**.

- Related Documentation**
- [Understanding vGW Series on page 3](#)
 - [Understanding the Multi-Center Feature on page 244](#)
 - [Understanding vGW Series Multi-Center Synchronized Objects on page 250](#)

Understanding the vGW Series Policy per vNIC Feature

This topic covers the vGW Series Policy per vNIC feature that allows you to configure separate firewall policies for individual interfaces, or virtual NICs (vNICs), configured on the same VM.

Before you use Policy per vNIC, you should be familiar with how to secure VMs and manage firewall policies, and you should have an overall understanding of the configuration of VMs that include more than one vNIC.

This topic includes the following sections:

- [About Policy per vNIC on page 258](#)
- [Why Use Policy per vNIC on page 259](#)
- [vNICs With Individual Policies and Smart Groups on page 259](#)
- [Viewing vNICs With Individual Policies on page 259](#)
- [Naming Conventions for vNICs on page 260](#)

About Policy per vNIC

You use the Settings module vGW Application Settings > Install Settings > Policy Per vNIC pane to enable the Policy per vNIC feature. You can enable the Policy per vNIC feature or you can allow the default capability that secures all vNICs on a VM in the same way. If you enable Policy per vNIC, you can still configure a policy for a VM that has only one vNIC.


If you do not enable Policy per vNIC, you can not configure individual policies for any vNICs on a VM that has more than one vNIC. In that case, all of the VM's vNICs inherit the same policy.

If you enable the Policy per vNIC feature, you can enable an option that allows you to exempt one or more vNICs on the same VM from requiring a firewall policy, effectively bypassing firewall security. When you enable this option, you can secure some individual vNICs with their own policies and leave other vNICs on the same VM unsecured.

You enable or disable Policy per vNIC at the global level: its configuration applies to all VMs that you secure using the same vGW Security Design VM. You cannot disable Policy per vNIC when individual vNICs have active policies applied to them.

You create policies for vNICs using the Firewall Manage Policy page. [Figure 125 on page 259](#) shows the policy page for the vNIC1 that belongs to the IT-WWW-DEV VM.

Figure 125: Policy for Single vNIC



Why Use Policy per vNIC

Policy per vNIC satisfies many requirements that emerge in a virtualized environment. For example:

- If your environment includes more than one PortGroup/vSwitch, you might want to have different policies for each of the networks that their VMs connect to.
- Your environment might include a server that connects both to the front end for customer interaction and to the back end for storage and management. You might want to disable the firewall on the back end but enforce it on the front end. In this case, you could use **Enable opt-out of firewalling per vNIC** feature.
- Your environment might include a single VM that has multiple vNICs attached to it, some of which have IPv6 addresses bound to them and some of which have IPv4 addresses bound to them. You can use the Policy per vNIC feature to apply different policy rules to vNICs passing IPv4 traffic from those used for IPv6 traffic, even when the vNICs are attached to the same VM. You could configure specific addresses for source or destination terms or you could use the predefined terms **Any-IPv4** and **Any-IPv6**.

vNICs With Individual Policies and Smart Groups

VMs for which the Policy per vNIC feature is used can be included in Smart Groups. You can choose whether membership in a Smart Group applies to the entire VM, that is, all of its interfaces, or only the vNICs that the Smart Group logic applies to. For example, an interface (a single vNIC) might belong to a port group or be connected to a certain VLAN which could qualify its membership in a Smart Group. For details on the relationship between vNICs and Smart Groups when Policy per vNIC is configured, see [“Understanding Policy per vNIC and Smart Groups for VMware Environments”](#) on page 266.


Viewing vNICs With Individual Policies

This section gives an overview of vNICs information as displayed by the vGW Security Design VM. See also [“Configuring and Displaying vGW Policies for Individual vNICs on the Same VM”](#) on page 263.

When the Policy per vNIC feature is enabled:

- vNICs are displayed under their VM in the VM Tree. The VM expands to show its individual vNICs. For example, as [Figure 126 on page 260](#) shows, IT-DB-Dev expands to show vNIC1 and vNIC2. Although this page shows the policy for the entire VM, you can select a vNIC and see only its policy.

Figure 126: VM with Multiple vNICs Shown in the VM Tree



- For operations that pertain to a VM, such as Introspection and Compliance, individual vNICs are not shown. They are treated in the same way as the VM that they belong to.

If a VM includes a vNIC that is not compliant, then the VM is considered noncompliant.

If you do not use the Policy per vNIC feature, the same policy is applied to all vNICs of a VM, and the VM is displayed as a single host in the VM Tree.



NOTE: When a vNIC with a policy is deleted, it no longer shows up in the list of vNICs with a policy. When you disable Policy per vNIC, the policies for all deleted vNICs are cleared. However these changes are not applied automatically. Consequently, if you create a vNIC again after having deleted it, the Apply Policy page for the VM might show that there are policy changes that have not been applied, but it would not state changes under Global, Group, and VM policies.

Naming Conventions for vNICs

vGW Series aligns with the convention for naming vNICs that is used by VMware in its vCenter:

- In VMware, naming of vNICs follows this convention: Network adapter 1, Network adapter 2, and so on. Numbering of vNICs begins with 1, not 0.
- In vGW Series, naming of vNICs follows this convention: VMx.nic1, VMx.nic2, and so on.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Configuring the vGW Series Policy per vNIC Feature on page 261](#)

Configuring the vGW Series Policy per vNIC Feature

This topic explains how to enable and configure the vGW Series Policy per vNIC feature that allows you to define separate policies for individual vNICs attached to the same virtual machine (VM).

Before you read this topic, read [“Understanding the vGW Series Policy per vNIC Feature” on page 258](#).



NOTE: For VMs that have multiple vNICs, you can still use the default configuration that allows you to use the same policy for all vNICs on your VMs. You are not required to use Policy per vNIC.

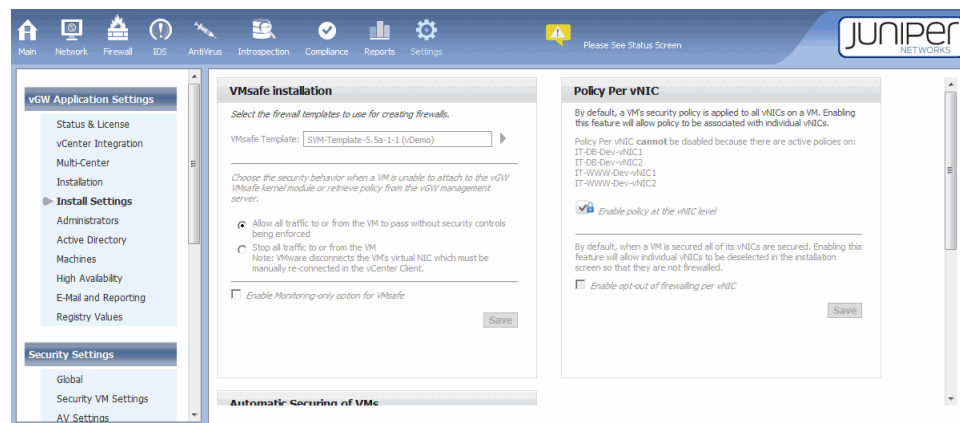
You can configure vNICs on the same VM to use:

- Separate policies for all vNICs on a VM.
- Separate policies on some vNICs on one VM while leaving other vNICs on the same VM unsecured.
- The same policy for all vNICs on a single VM (default).

You cannot disable Policy per vNIC when individual vNICs have active policies applied to them.

[Figure 127 on page 261](#) shows the Install Settings page that you use to enable vGW Policy per vNIC and define its behavior.

Figure 127: Policy Per vNIC



To enable Policy per vNIC:

1. In the Settings module vGW Application Settings section, select **Install Settings**.
2. To enable the feature globally, in the Policy Per vNIC pane, select the **Enable policy at the vNIC level** check box.

3. Optionally, select the **Enable opt-out of firewalling per vNIC** check box if you want to secure some vNICs but not others on the same VM. See [“Configuring Policy per vNIC to Secure Only Some of a VM’s vNICs” on page 263](#).

When new interfaces are added to a VM that includes vNICs that are not secured, the new vNICs are automatically secured. If you want them not to be secured, you must manually unsecure them. The following procedure explains how to remove security from a vNIC.

If you disconnect a vNIC from a port group, that is, un-selected it, the vNIC becomes unsecured. A warning message on the Installer dialog shows the state of the vNICs.



CAUTION: If you select “Enable opt-out of firewalling per vNIC” on the Policy Per vNIC pane, vNICs can not be secured individually if they belong to the same port group.

This procedure explains how to remove a security policy from a vNIC, that is, *unsecure* it. To unsecure a vNIC:

1. Select the vGW Security Design VM Settings module.
2. In the vGW Application Settings section, select **Installation**.
3. Before you unsecure the vNIC, delete any policies applied to it.
4. In the Secured Network pane, select the vNIC that you want to leave unsecured, and click the **Unsecure** arrow.

The vGW Security Design VM presents a message that asks you whether you want to unsecure the vNIC or the entire VM.

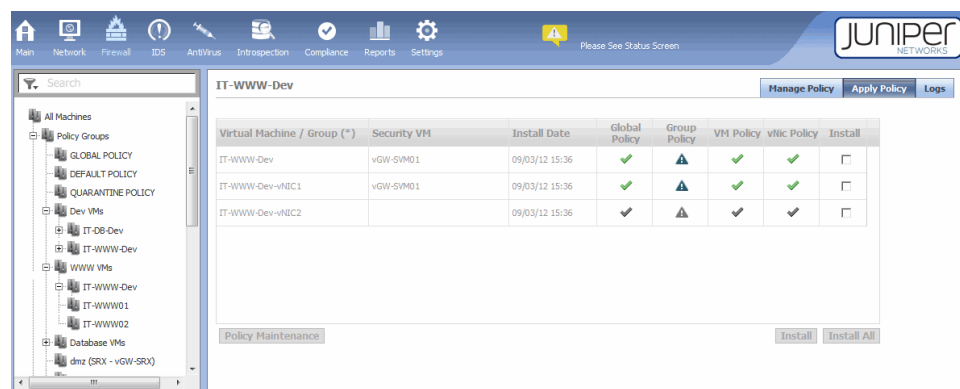
If you add a new vNIC to a VM that contains vNICs that are not secured, the new vNIC is automatically secured. If you want to unsecure it, you must do it manually as explained previously.

You use the Firewall module pages to create and apply policies for vNICs that belong to a VM with multiple vNICs and for which you use the Policy per vNIC feature.

[Figure 128 on page 263](#) shows the Firewall module Apply Policy page for the IT-WWW-DEV VM with multiple vNICs. To apply the policies, you must select the **Install** check box and click **Install (Install All)**.

For details on how to define individual policies for vNICs, see [“Configuring and Displaying vGW Policies for Individual vNICs on the Same VM” on page 263](#).

Figure 128: Applying Policy to Individual vNICs



Related Documentation

- [Understanding vGW Series on page 3](#)
- [Configuring Policy per vNIC to Secure Only Some of a VM's vNICs on page 263](#)

Configuring Policy per vNIC to Secure Only Some of a VM's vNICs

The Policy per vNIC feature includes an option that allows you to secure some of your vNICs and leave others unsecured. To use this option, you must enable Policy per vNIC. You use the Policy per vNIC pane on the Install Settings page to enable Policy per vNIC and to select the Enable opt-out of firewalling per vNIC option.

If you select the Enable opt-out of firewalling per vNIC option, the unit of configuration is the VM and port group. That is, vNICs cannot be secured individually if they belong to the same port group. This behavior protects against your having a secured and an unsecured connection to the same port group.



NOTE: When new interfaces are added to a VM that includes vNICs that are not secured, the new vNICs are automatically secured. If you want them not to be secured, you must manually unsecure them.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Configuring and Displaying vGW Policies for Individual vNICs on the Same VM on page 263](#)

Configuring and Displaying vGW Policies for Individual vNICs on the Same VM

This topic covers how to configure policy rules for individual vNICs that belong to the same virtual machine (VM) when the Policy per vNIC feature is enabled. It also explains how vNICs are displayed in the VM Tree. You use the Firewall module of the vGW Security Design VM to configure and apply policies to vNICs.

When Policy per vNIC is enabled and multiple vNICs for the same VM have been configured they are presented in the VM Tree nested beneath the VM that they belong to.

The VM Tree displays the state of a vNIC in the following ways:

- The VM Tree displays the state of a vNIC in the following way:
 - A disabled vNIC is shown with an icon that indicates that traffic on the vNIC is not protected by the vGW Security VM firewall.
 - If a VM contains vNICs with individual firewall policies and the VM belongs to a group, the vNICs that are members of the group are shown as active. The vNICs that do not belong to the group are shown, but they are grayed out indicating that they are not part of the group.



NOTE: vNIC numbers can change when one vNIC is deleted. For example, if a VM contains vNIC1 and vNIC2 and you remove vNIC1, then vNIC2 becomes vNIC1. If you have manually created policies for both vNIC1 and vNIC2, the enforced policy is also changed so that the correct policy for the vNIC remains with it.

- vNICs are displayed in the VM Tree:
 - If more than one vNIC is configured for a VM.
 - When there remains one vNIC configured for a VM and a policy is applied to it.

In this case, originally there were multiple vNICs configured for the VM, each with its own policy, and all except one of them was deleted. It is still possible for you to edit or delete the policy on the remaining vNIC.

vNIC policies are shown above the policy for the VM that they belong to, in the order in which they were defined.

- vNIC policies are enforced after the policy for the VM that they belong to.
- When you select a VM in the VM Tree, policies for the vNICs that belong to it are shown as read-only.
- When you select a vNIC in the VM Tree, the policy for that vNIC is shown, and you can edit it. All other policies are shown as read-only. Policies for other vNICs are not displayed.
- From the perspective of the rule base, vNIC policies behave in the same way as other policy types:
 - If the vNIC is selected in the VM Tree, the policy for it can be edited. If the VM is selected, the vNIC policies are greyed out indicating that they can not be edited.
 - For unsecured vNICs, the vNIC header is shown. Instead of rule information, the following message is displayed: "This interface is configured to bypass firewall enforcement".

When Policy per vNIC is enabled, the Apply Policy table reflects the vNIC configuration in the following way:

- vNICs are displayed as rows in the Apply Policy table. When Policy per vNIC is disabled or a VM does not contain multiple vNICs, the table displays information for the VM as usual. See [“Understanding the vGW Series Firewall Module” on page 93](#) for details on the Apply Policy table.
- When you select the VM in the VM Tree, the policy for it is displayed. However, there is a table entry for each vNIC, but it reads “(no rules)”.
- Each vNIC has its own policy. If all vNICs except one are removed from the VM, the remaining vNIC is displayed in the table. Its policy can be edited or deleted.



NOTE: You can not disable the Policy per vNIC feature if there are policies configured for any vNIC or groups containing the vNIC. You must first delete the policies.

You use the Firewall module Manage Policy tab to add rules for individual vNICs in the same way that you configure other policy rules.

This procedure explains how to define policies for the following example. For additional details on how firewall policy rules are configured, see [“Understanding the vGW Series Firewall Module” on page 93](#).

This example assumes that the Policy per vNIC feature is enabled. For details on how to enable Policy per vNIC, see [“Configuring the vGW Series Policy per vNIC Feature” on page 261](#). The example assumes that the administrator wants to configure separate policies for each of the following three vNICs on a VM called MIS-Fileserver that is used as a file server:

- vNIC1 (MIS-Fileserver-vNIC1) is dedicated to network connections, and it requires a policy whose protocol specification allows https and ssh traffic.
- vNIC2 (MIS-Fileserver-vNIC2) whose policy allows the iSCSI protocol to use to link data storage facilities.
- vNIC3 (MIS-Fileserver-vNIC3) that is used for management that allows SNMP protocol traffic.

To configure policies for these vNICs:

1. In the vGW Security Design VM, select the Firewall module.
2. In the VM Tree, locate the MIS-Fileserver VM, and expand it to display the vNICs.
3. Select vNIC1.

When you select the vNIC, the policy page for it is displayed. The policy is called vNIC Policy for MIS-Fileserver-vNIC1.

4. Beneath the Global Policy line is a line labeled “vNIC Policy for MIS-Fileserver-vNIC1” that allows space for you to enter a policy rule for the vNIC.

Click **Add**.

5. In the Sources column for the rule, leave **Any**.
6. In the Protocols column for the rule, click **Any** to display a list of protocols.
 - a. In the Filter box, enter **https**. The list is scrolled to https (443/tcp). Select it and click the right-facing **Arrow** to move it to the Selected Protocols box.
 - b. In the Filter box, enter **ssh**. The list is scrolled to ssh(22/tcp). Select it and click the right-facing **Arrow** to move it to the Selected Protocols box.

Click **Save**.

When Policy per vNIC is enabled, the Apply Policy table contains an additional column to indicate policy state for the vNIC.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding Policy per vNIC and Smart Groups for VMware Environments on page 266](#)

Understanding Policy per vNIC and Smart Groups for VMware Environments

You use the vGW Security Design VM Settings module vGW Application Settings > Install Settings > Policy Per vNIC pane to enable the Policy per vNIC feature. When it is enabled, you can add individual vNICs to a Smart Group. When you configure a Smart Group, you can specify whether requirements for membership in the group apply to an entire VM, that is, all of its interfaces, or only to the vNICs that the logic pertains. For example, Smart Group criteria might specify that the vNIC must belong to a port group or that it must be attached to a VLAN to gain membership in the group.

The ability to configure Smart Groups for vNICs is available only when Policy per vNIC is enabled. You can configure this information when **Advanced Attributes** is selected.

After you configure the group, you can test it. When you click **Test**, the results show the vNIC extensions, not just the VM name.

You can use the following Smart Group attributes to configure groups to include vNICs. These attributes do not pertain to the VM as a whole.

Table 16: Smart Group Attributes for vNICs When Policy per vNIC Is Enabled

Smart Group Attribute Definition	Data Type	Comment
vf.firewall	String	Is this VM a vGW Security VM?
vf.group	Multi String	Comma-separated string of all vGW groups to which a VM belongs.
vf.has_installed_group_policy	Boolean	Does the VM have a non-default group policy installed?

Table 16: Smart Group Attributes for vNICs When Policy per vNIC Is Enabled (*continued*)

Smart Group Attribute Definition	Data Type	Comment
vf.has_installed_policy	Boolean	Does the VM have an installed security policy?
vf.monitored	Boolean	Is the VM currently being monitored by the vGW Security Design VM?
vf.secured	Boolean	Is a VM currently secured by the vGW Security Design VM?
vf.secured_active	Boolean	Is the VM actively protected by vGW?
vi.host.vmkernl.isolated.vlan	Boolean Value	Is the vmkernel management network on this hypervisor on an isolated VLAN?
vi.host.vmkernl.isolated.vswitch	Boolean Value	Is the vmkernel management network on this hypervisor on an isolated vSwitch?
vi.ipv4	IPv4 (multi value)	The IP addresses as known on a VM.
vi.ipv6	IPv6 (multi value)	<p>The IP addresses as known on a VM. They can be coded as single addresses or an address range.</p> <p>Example Addresses:</p> <ul style="list-style-type: none"> • 2001:0db8:5a300000:0000:0a2e:0370:7334 • fe80::202:b3ff:fe1e:8329
vi.pg_security.forgedtransmits	Boolean Value	Is VM connected to a port group which allows forged MAC addresses (MACs other than defined in the VMX)?
vi.pg_security.macchanges	Boolean Value	Is VM connected to a port group which allows reception of unknown MAC addresses (MACs other than defined in the VMX)?
vi.pg_security.promiscuous	Boolean Value	Is VM connected to a promiscuous port group?
vi.portgroup	String Value	Port groups on the virtual switch this VM is actively connected to. Port Groups for disconnected vNICs will not be included. (For a running/suspended VMs this will be the port groups actually connected. For a stopped VM, this value is the port groups that are connected at power-on.)

Table 16: Smart Group Attributes for vNICs When Policy per vNIC Is Enabled (*continued*)

Smart Group Attribute Definition	Data Type	Comment
vi.portgroup.all	String Value	Port groups on the virtual switch this VM configured to be connected to, this list includes port groups even if the vNIC is disconnected. (For a running/suspended VMs this will be the port groups actually connected. For a stopped VM, this value is the port groups that are connected at power-on.)
vi.pvlan	Numeric Value	Private VLAN values for connected port groups.
vi.pvlan.all	Numeric Value	List of all Private VLANs in use by this VM, includes vNICs in both connected and disconnected states.
vi.vlan	Multi-value integer	VLANs of connected port groups.
vi.vlan.all	Multi-value integer	VLANs of all interfaces.
vi.vmsafe_configured	Boolean	Is VMsafe firewall security enabled for this VM?
vi.vmsafe_dvfilter	Multi String	The dvfilters protecting this VM.
vi.vmsafe_initfailmode	Enumeration	If VMsafe is unable to initialize, what is the network connectivity choice for this VM?
vi.vnic.count	Numeric Value	Number of connected vnics.
vi.vswitch	Multi String	vSwitch VM is connected to.

You use the attributes shown in [Table 16 on page 266](#) to define a Smart Group. The Smart Group editor has two modes: basic and advanced. Basic mode lets you select one to many attributes and assign an All or Any constraint. You simply add rules by clicking the + sign. Advanced mode allows you to configure the Smart Group for vNICs.

1. In the Security Settings section of the vGW Security Design VM Settings module, select the **Groups** subsection.
2. Click **Add Smart Group** on the displayed page.
3. Click **Advanced** at the top of the page to display vNIC group options.
4. In the Add Group definition pane, enter a name for the Smart Group. For this example, enter **Apache Web Servers**.
5. Click **Enable vNIC membership** to specify that group membership pertains to vNICs, and not the VM.

6. Select the **All** option button in the Matches section.
7. Click the down arrow to display a list of attributes. Select the attribute **vi.name**, select **Contains**, and enter **www**.
8. Click the **+** mark at the end of the row to display another row.
9. Select the attribute **vf.application**, select **Contains**, and enter **www**.
10. Under Group Attributes, select **Policy Group** allow a policy to be associated with this group.
11. Select **Medium** as the Priority level, and assign it a precedence of **2** in the Precedence within Level.
12. Select Manual.

This allows you to use the Settings module and apply a policy to the group using the Firewall Apply Policy tab.

1. Specify a name for the group and configure its attributes.
2. Click **Enable vNIC membership** to specify that group membership pertains to vNICs, and not the VM.
3. Click **Test** to view the results of your configuration.

The test results show the VM name with the vNIC extension that the Smart Group logic applies to.

When you view a Smart Group in the VM Tree and display the VM with its nested vNICs, vNICs that belong to the group—that satisfy the group's logic criteria—are displayed as usual. vNICs that do not belong to the group are greyed out.

Whether a vNIC is secured or not is indicated as usual for all of the VM's vNICs.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Configuring and Displaying vGW Policies for Individual vNICs on the Same VM on page 263](#)

Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module

This topic includes the following sections:

- [Configuring an Administrator Account on page 269](#)
- [Changing Administrator Passwords on page 272](#)

Configuring an Administrator Account

Different categories of IT staff members may need to access the vGW Security Design VM interface for various purposes. For example, network engineers can take advantage of the network statistics charts and information on connections, top protocols used, top sources, and top destinations. Security engineers can use the Firewall module to design

and apply policies for VMs and the Settings module's vGW Application Settings > Installation page to deploy vGW Security VMs to ESX/ESXi hosts to secure them.

[Table 17 on page 270](#) defines the built-in user types that vGW Series provides to accommodate common roles and requirements, and it describes their privileges.

Table 17: vGW Series Built-In Administrator User Types

Global Admin	<p>This administrator has the highest level of system privileges, including the ability to create accounts for additional administrators.</p> <p>The global administrator has many privileges including the ability:</p> <ul style="list-style-type: none"> • to create firewall policies and install firewalls (vGW Security VMs) on ESX/ESXi hosts to be secured. • configure features such AntiVirus, IDS, and VM Introspection Compliance for VMs. • select port groups and VMs for insertion in and removal from a secured network. <p>This administrator can also change his own password and reset the passwords of other administrators. Having the ability to reset the password for another administrator is useful when an administrator forgets his password. For details see “Changing Administrator Passwords” on page 272.</p>
VM Admin	<p>These administrators have many privileges, including the ability to:</p> <ul style="list-style-type: none"> • modify policies and settings configurations. <p>The administrator is allowed to change firewall security policies, including IDS.</p> <ul style="list-style-type: none"> • configure AntiVirus and VM Introspection Compliance. • configure mirroring of inter-vm traffic, the ability to configure rules that specify external inspection devices. <p>Additionally, the global administrator can grant VM Admins “Install Firewall Policy” privilege. This privilege allows a VM Admin to distribute a policy after it has been changed and saved by any administrator who has the privilege to modify security policies.</p>
Network Monitoring	<p>These administrators can view:</p> <ul style="list-style-type: none"> • all network-related pages, for example pages that show statistics and graphs. • all tabs of the Main module, including Status and Events and Alerts, and Logs. <p>These administrators are not allowed to modify any Settings pages, but they can view IDS Alerts, if IDS is configured, view AntiVirus scans, and they can view but not modify VM Introspection and Compliance results.</p>

To create an administrator account:

1. From Settings module vGW Application Settings > Administrators page, click **Add**.

Figure 129 on page 271 shows the Administrators page > Add Administrator pane that you use to define permissions for a new administrator and add the administrator to the system.

This example configuration specifies that authentication is performed internally by vGW Series, not by Active Directory (AD), which could also be used. In this example, the VM Admin admin-security-example administrator is allowed to modify policy and settings and push firewall policies to vGW Security VMs.

Figure 129: Creating a VM Admin Administrator Account

The screenshot shows the Juniper vGW Series web interface. On the left is a navigation menu with categories like vGW Application Settings, Security Settings, and Appliance Settings. The 'Administrators' page is active, displaying a table of existing administrators. Below the table is the 'Add Administrator' dialog box. In this dialog, the 'Authentication Type' is set to 'Internal', the 'Username' is 'admin-security-example', and the 'Full Name' is 'admin-security-example'. The 'Type' is set to 'VM Admin'. Under 'Permissions', several checkboxes are selected: 'Modify policy and settings', 'Allow monitoring of inter-VM traffic', 'Allow scheduling and starting VM introspection scans', 'Allow configuration backup and restore', and 'Install Firewall policy'. At the bottom, there is a 'Change password' section with fields for 'New Password' and 'Confirm Password'.

Username	Full Name	Type	Authentication Type
admin	Default Global Admin	Global Admin, Console administrator	Internal

2. In the **Authentication Type**: area, select the button associated with the kind of authentication to be used for this administrator. You can use Active Directory (AD) as a means of authentication rather than storing the credentials locally. In this case, Active Directory must first be enabled through the Settings module > vGW Application Settings > Active Directory page. For details on AD authentication, see [“Setting Up Active Directory for vGW Series Administrator Authentication”](#) on page 275.
3. In the **Username**: and **Full Name**: fields, enter the user names for the administrator.
4. In the **Type**: area, select the button associated with the type of administrator account that you want to create. See [Table 17 on page 270](#).
5. In the **Permissions**: area select the permissions that you want to grant to the administrator. Notice that for VM Admin you can select “Modify policy and settings” and “Install Firewall policy”, but if you select Network Monitoring you cannot select any of these permissions. See [Table 17 on page 270](#) for allowed permissions.
6. Specify a password and confirm the password.
7. Click **Save**.

After you save the configuration, the administrator definition is added to the Administrators table, as shown in [Figure 130 on page 272](#).

Figure 130: Adding a New Administrator

Administrators				Filter <input type="text"/>
Username	Full Name	Type	Authentication Type	
admin	Default Global Admin	Global Admin, Console administrator	Internal	
admin-security-examp	admin-security-example	VM Admin	Internal	



NOTE: At any time, you can click the table row for an administrator definition to display the Edit Administrator pane that shows the configuration. From the Edit Administrator pane you can modify the permissions and password and save the modified definition.

Changing Administrator Passwords

Whether you are a global administrator (Global Admin), an administrator whose account is defined as a VM Admin, or an administrator with Network Monitoring permissions, you can use the Settings module vGW Application Settings > Administrators page to change your password.

This section includes the following sections that explain the simple process and requirements:

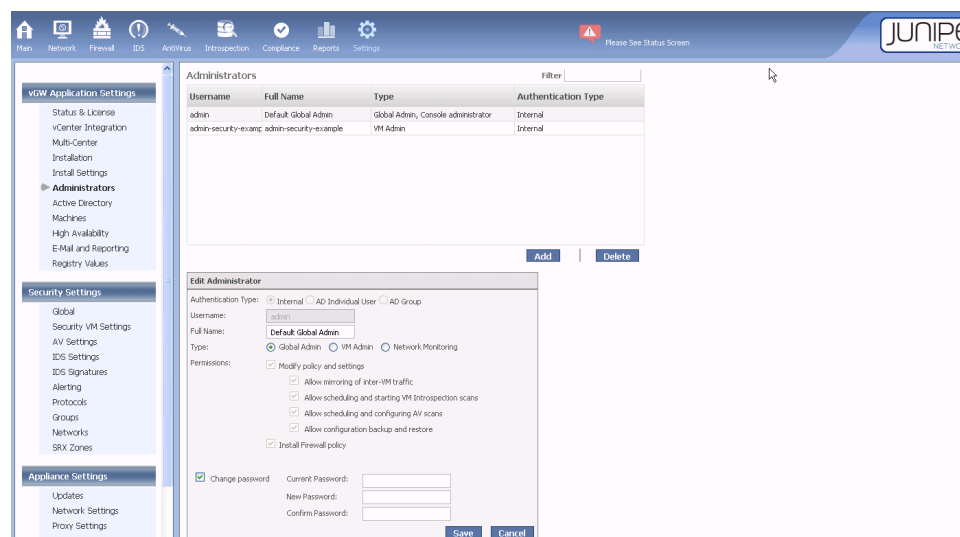
- [Global Administrator: Changing Your Own Password on page 272](#)
- [Global Administrator: Changing the Password of Another Administrator on page 273](#)
- [VM Administrator and Network Monitoring Administrator Accounts: Changing Your Own Password on page 274](#)

Global Administrator: Changing Your Own Password

As the global administrator (Global Admin), when you select your own row in the Administrators table, the **Edit Administrator** dialog box appears showing the configuration for your account. To change your own password, you must first enter your current password followed by the new one.

When you select the **Change password** check box, the **Current Password:** and **New Password:** boxes appear, allowing you to change your password. You must also enter the new password in the **Confirm Password:** box. After you enter the new password, click **Save**. [Figure 131 on page 273](#) shows this dialog box.

Figure 131: Changing the Global Administrator Password



Global Administrator: Changing the Password of Another Administrator

When you want to change the password of another administrator—such as an administrator whose account is defined as a VM Admin or for an administrator with Network Monitoring permissions—you are not required to enter that administrator's current password. Not having to enter the current password for another administrator allows you to provide that administrator with a new password when they forget their current one.

As the global administrator, when you select the row for another administrator in the Administrators table, the **Edit Administrator** dialog box appears, showing the configuration for that administrator's account.

As [Figure 132 on page 274](#) shows, when you select the **Change password** check box, the **New Password:** and **Confirm Password:** boxes appear, allowing you to change the password for the administrator whose account configuration is displayed. After you enter the new password, click **Save**.

Figure 132: Global Administrator Changing the Password of Another Administrator

Edit Administrator

Authentication Type: ☒ Internal ☐ AD Individual User ☐ AD Group

Username:

Full Name:

Type: ☐ Global Admin ☒ VM Admin ☐ Network Monitoring

Permissions:

- ☒ Modify policy and settings
 - ☒ Allow mirroring of inter-VM traffic
 - ☒ Allow scheduling and starting VM Introspection scans
 - ☒ Allow scheduling and configuring AV scans
 - ☒ Allow configuration backup and restore
- ☐ Install Firewall policy

☒ Change password New Password: Confirm Password:

Save **Cancel**

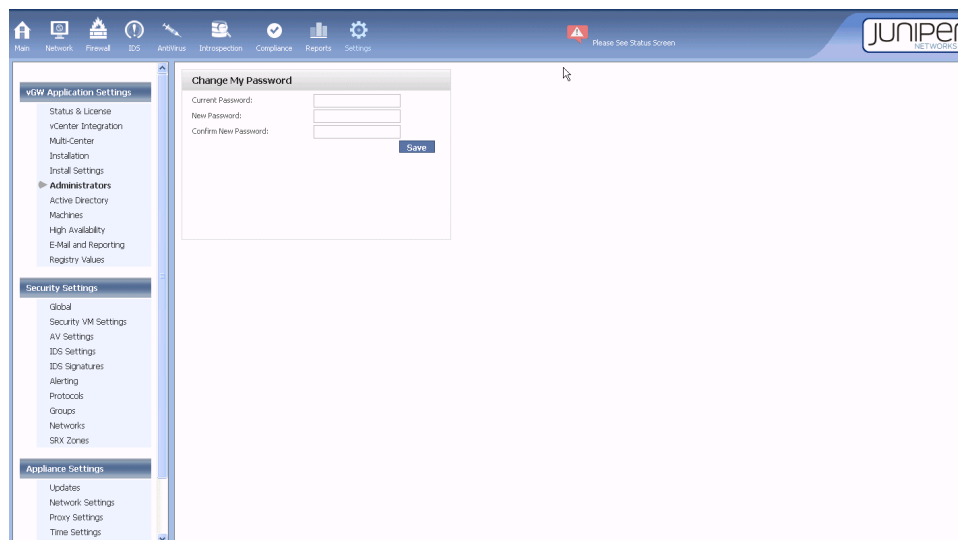
VM Administrator and Network Monitoring Administrator Accounts: Changing Your Own Password

After the global administrator (Global Admin) defines an administrator account for you, you can change the password that was specified during the configuration. In this case, the global administrator conveys the password to you. You can also change your password at any time after you change it initially.

When you select Administrators, the change password dialog box appears. To change your password, you must first enter your current password followed by the new one.

To change your password, enter your current password in the **Current Password:** box and your new password in the **New Password:** box. You must also enter the new password in the **Confirm Password:** box. Then click **Save**. See [Figure 133 on page 275](#)

Figure 133: Administrators Changing Their Password



Related Documentation

- [Understanding vGW Series on page 3](#)
- [Setting Up Active Directory for vGW Series Administrator Authentication on page 275](#)
- [Configuring vGW Series Firewall Policies on page 114](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Security Design VM on page 67](#)

Setting Up Active Directory for vGW Series Administrator Authentication

This topic covers use of Active Directory (AD) for administrator authentication. First it explains how to enable AD support for vGW Series, which you must do before you can configure administrator authentication to use it. Then it explains how to configure it as the authentication type for an administrator.

You can use AD with vGW Series for administrator authentication instead of storing the authentication information locally in the vGW Security Design VM database. vGW Series supports AD over IPv4 and IPv6 networks.

Administrators can use their AD credentials to log in to the vGW Security Design VM. vGW Series checks AD for the credentials, and, based on the settings, it allows the user to log in to vGW Security Design VM or it denies the user access.

To set up the vGW Series to work with AD:

1. Define the Name (or IP address) of the AD server on the Active Directory configuration page.
2. Set the appropriate port. By default, port TCP 636 (LDAPS) is used. However, you can use 389 LDAP+STARTTLS or configure a custom port.

Enable your network to give the vGW Security Design VM access to this port to the server.

3. After you select the name or IP address, port, and default search base, select **Test** or **Save** to view the fingerprint used to validate the communication destination and to initiate all future communication through encryption.

When you select **AD Group** for **Authentication type**, a dialog box is displayed allowing you to enter the user ID and password to use to log in to AD to get the group list.



NOTE: AD must be enabled for you to select **AD Group** as the authentication method. Use the Settings module vGW Application Settings > Active Directory page to enable it, as described previously.

If there are more than 100 configurable groups, vGW Series presents the following alert message:

“There are too many groups in Active Directory to be displayed in a drop-down list. Please fill in the name of the AD Group.”

Rather than displaying a drop-down list of group names, the AD Group Name field is presented as a text box in which you can enter the name of the group.

When you save the configuration, vGW Series checks AD to ensure that the group exists, based on the name that you entered. If the group does not exist, vGW Series displays the following message:

“The AD Group *name* does not exist in Active Directory.”

To create users or groups to be authenticated through the configured server lookup process:

1. Select the Settings module > vGW Application Settings > Administrators page.
2. Add administrators. Set the authentication type to **AD Individual User** or **AD Group**.
 - For AD Individual User, the account is authenticated with AD credentials and all privileges are applied according to defined vGW Series settings.
 - For AD Group, the name of an existing group in AD is used and privileges are assigned to it. The AD lookup is used to authenticate the user to determine that he is a member of the group. If so, he is granted access to vGW Series.

Related Documentation

- [Adding New vGW Series Administrator Definitions, Permissions, and Authentication Using the Settings Module on page 269](#)
- [Understanding vGW Series on page 3](#)
- [Configuring vGW Series Firewall Policies on page 114](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Security Design VM on page 67](#)

- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)

Adding and Editing vGW Series Machines Definitions (VMware)

This topic covers the Machines page that you use to define IP addresses and other information for new machines. These machines include both VMware ESX/ESXi hosts and virtual machines (VMs) that you define for your environment. You also use this page to view or edit information about machines that are already defined, including those that are discovered automatically. Machines can have IPv4 or IPv6 addresses.

This topic describes a new parameter provided with vGW Series Release 5.5—Log Tags—that allows you to specify tags that are added to syslog output. You can use these tags to sort on syslog feed.

This topic includes the following sections:

- [Adding a Machine on page 277](#)
- [Viewing Machine Information on page 279](#)

Adding a Machine

Normally the IP address for a machine is “auto-discovered”, obtained through VMware Tools. For systems without VMware Tools, you can use the Settings module Applications Settings > Machines page to manually add addressing information for a machine. You can also specify additional information for a machine, such as Log Tags and Smart Tags.

To configure information for a machine, enter:

- **Name:** This is the name of the machine (VM). By default, it is set to synchronize with the VMware vCenter. However, you can detach it by clearing the Synchronize name with vCenter checkbox.

In Release 5.5, vGW Series adds the name to syslog output, instead of adding just VM_ID. The name is relative to the VM that is either the source (src) or destination (dst) of the log flow. For example, dst_name="mini-5-1" or src_name="mini-5-1". hr".

- **Description:** Give a brief description of the machine.
- For the machine's address, use DNS for the machine name or enter its address explicitly.
 - **DNS name:**

If you define a machine in this section, it is identified in the network tables by its name rather than by its IP address.

 - Specify the machine's DNS name.
 - To obtain the name through a DNS query, click **Query via DNS**.
 - **IP Address:** Specify the machine's IP address explicitly.
- **Smart Tags:** Optionally, configure Smart Tags to assign identifiers to the machine that can be used for VM Smart Groups or policy creation.

The syntax for a Smart Tag is attribute-value. You can define multiple tags separated by semicolons, for example: `finance;pci=true;audited=true`.

- **Log Tags:** Optionally, specify Log Tags to be added to Syslog entries for this machine.

In conjunction with the VM name that is added to Syslogs as of vGW Series Release 5.5, this option allows you to specify any tag that you want to use to be added to the syslogs. For example, you can use this tag to associate certain VMs with a Tenant or Department such as customer A's VMs are tagged with 'cust-a' and which are then sorted automatically from the syslog feed by parsing on this tag.

Similar to the VM name tag described above, these tags are relevant on direction of the flow (`src_log_tag` or `dst_log_tag`). For example, `dst_log_tag="testLogTags"`. These logs are issued when vGW Series processes secured VM traffic or files.

The Log Tag string pertains to this VM only. It cannot exceed 200 characters. Figure [Figure 134 on page 279](#) shows the edit screen that includes a log tag for a machine that was already added. Figure [135 on page 279](#) shows the resulting syslog entry.

The Log Tag string pertains to this VM only. It cannot exceed 200 characters.

- **Type:** This is the type of machine, for example, ESX/ESXi server, external machine
- **Monitoring Groups:** Monitoring groups that the machine belongs to.
- **Policy Groups:** Policy Groups that the machine belongs to.
- **VMSafe Protected:** Whether the machine is secured by vGW Series.

When you select a VM, as opposed to an ESX/ESXi server, and display the Edit Machine box for it, this information is displayed for it.



NOTE: If you click **Advanced. . .** You can change the behavior if vGW Series fails to connect to the kernel (`failopen` or `failclosed`).

You can also use the Machines page to edit information for an existing machine. See [Figure 134 on page 279](#).

Figure 134: Configuring Machines Information

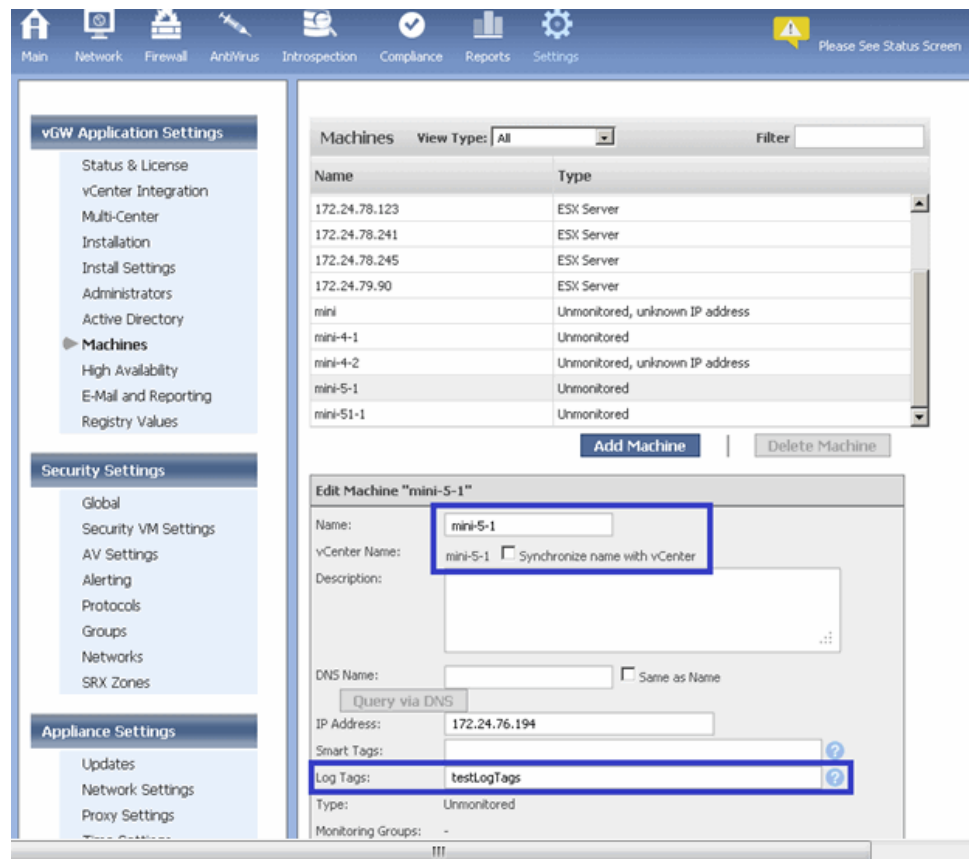


Figure 135: Syslog Entry Including VM Name and Log Tag

```
action=allow vm_id=0 ip_proto=udp rule=7 type=fw src_id=0 dst_id=29 svm_id=32 src_name= dst_name="mini-5-1" src_log_tag= dst_log_tag="testLogTags"
```

Viewing Machine Information

You can view information about machines that are already defined. You can use the **View Type:** box to sort the list by machine type. You can sort by ESX servers, external machines, monitored, unmonitored, and secured machines.

You can use the Filter box to search by a portion of an IP address or machine name or type.

- Related Documentation**
- [Understanding vGW Series on page 3](#)
 - [vGW Security Design VM Modules \(VMware\) on page 6](#)

Configuring vGW Series E-Mail and Reporting Applications Settings

You use the Settings module vGW Application Settings E-Mail and Reporting section to configure the e-mail server and account information. This information is used throughout vGW Series for distributing status and log messages and reports.

During the installation of the vGW Security Design VM, you can configure the parameters required to generate automated reports.

To configure or change these parameters, use the Settings module > vGW Application Settings > E-Mail and Reporting page, and enter the new values.

The following list describes the e-mail settings and configuration parameters:

SMTP Server—Hostname or IP address where e-mail is sent. You can specify either a valid IPv4 address or IPv6 address.

SMTP Port—Port used by the mail server (common values are 25 or 465 for encrypted).

Authenticate—If authentication to the mail server is required, check this option.

TLS Authenticate—If the mail server uses TLS encryption, select this option.

SMTP—If authentication is required, use this user account.

E-mail From—Text that appears in the From field in e-mail messages.

E-mail To—Text that appears in the To field in e-mail messages.



TIP: You can troubleshoot mail server configuration errors by clicking **Test Mail Server** before saving parameter changes.

The reporting module settings configuration parameters are:

Default e-mail From—Text that will appear in the From field of e-mail messages by default.

Mail Subject —Text you want inserted in the Subject line for messages sent by the Reporting module.

Mail Content —Text you want inserted in the content section of messages. (The report itself is attached as a PDF file.)

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security VM on page 68](#)

CHAPTER 16

vGW Series Security Settings

- [Understanding the vGW Series Security Settings on page 281](#)
- [Configuring Global Settings Using the vGW Series Settings Module \(VMware\) on page 282](#)
- [Understanding the vGW Security VM Settings on page 286](#)
- [Understanding and Configuring the vGW Series AntiVirus Settings on page 291](#)
- [Understanding and Configuring IDS Settings on page 292](#)
- [Understanding and Configuring IDS Signatures Settings on page 295](#)
- [Understanding the vGW Series Security Alert Settings on page 297](#)
- [Understanding vGW Series Protocols Support on page 299](#)
- [Understanding vGW Series Groups on page 300](#)
- [Automatically Applying Policy Rules to VMs in Policy Groups on page 303](#)
- [Understanding vGW Series Smart Groups on page 305](#)
- [About Using vGW Series Attributes for VMware on page 306](#)
- [Creating vGW Series Smart Groups for VMware on page 307](#)
- [vGW Series Attributes for VMware on page 311](#)
- [Understanding the Settings Module Networks Settings on page 316](#)
- [Understanding the vGW Series SRX Zones Settings on page 317](#)

Understanding the vGW Series Security Settings

The Settings module Security Settings section controls the core vGW Series functions. The Security Settings module brings together configuration of many parts of your deployment so that you can configure or change them in one place.

The Settings Module > Security Settings includes the following sections:

- Global
[See “Configuring Global Settings Using the vGW Series Settings Module \(VMware\)” on page 282.](#)
- Security VM Settings
[See “Understanding the vGW Security VM Settings” on page 286.](#)
- AV Settings

See [“Understanding and Configuring the vGW Series AntiVirus Settings”](#) on page 291

- IDS Configuration

See [“Understanding and Configuring IDS Settings”](#) on page 292.

- IDS Signatures

See [“Understanding and Configuring IDS Settings”](#) on page 292.

- Alerting

See [“Understanding the vGW Series Security Alert Settings”](#) on page 297.

- Protocols

See [“Understanding vGW Series Protocols Support”](#) on page 299

- Groups

See [“Understanding vGW Series Groups”](#) on page 300.

- Networks

See [“Understanding the Settings Module Networks Settings”](#) on page 316.

- SRX Zones

See [“vGW Series and SRX Series Security Zones”](#) on page 359.

**Related
Documentation**

- [Understanding vGW Series](#) on page 3
- [Understanding the vGW Security Design VM](#) on page 67

Configuring Global Settings Using the vGW Series Settings Module (VMware)

This topic covers vGW Series Global settings. This topic includes the following sections:

- [vGW Series Global Settings Overview](#) on page 282
- [Global Settings](#) on page 283
- [vGW Series IPv6 Support and Global Settings](#) on page 285

vGW Series Global Settings Overview

The Settings module Security Settings > Global page allows you to identify the external inspection devices to send traffic to for further analysis, a Syslog server to use for external logging, global rules, and NetFlow configuration information that identifies where to send connection flow data. It also allows you to specify whether to allow or drop certain types of traffic, such as non-IP traffic.

Global settings that you configure apply to all vGW Security VMs *unless* you configure different information for a particular vGW Security VM. For details on using the Settings > Security Settings > Security VM Settings page to override the global configuration for individual vGW Security VMs, see [“Understanding the vGW Security VM Settings”](#) on page 286.

Figure 136 on page 283 shows the Global page.

Figure 136: Global Settings Page

The screenshot shows the Juniper vGW Series Global Settings page. The left sidebar contains navigation links for vGW Application Settings and Security Settings. The main content area is divided into several panes:

- External Inspection Devices:** A table with columns for Name and IP Address, and a Save button.
- External Logging:** Radio buttons for No Syslog, Send Syslog from vGW management server, and Send Syslog from Firewalls. A checkbox for Send firewall logs to vGW management server is checked. Fields for Syslog Server and Syslog Server Port (514) are present, along with a Save button.
- Global Settings Rules:** A table with columns for Rule, Allow, and a Save button. Two rules are listed: Rule 1 (IPv6 traffic) and Rule 2 (Non-IP and non-ARP traffic).
- NetFlow Configuration:** A checkbox for Enable, fields for NetFlow collector address and NetFlow collector port (2555), and a Save button.

Global Settings

The Global settings page contains the following panes:

- **External Inspection Devices**—This pane allows you to enter the names and IP addresses of devices to which traffic can be sent for further analysis, such as Intrusion Detection Systems and Network Analyzers. External inspection device must be capable of terminating a GRE tunnel. They must be reachable/routable from the vGW Security VM to every one of the external devices.

The vGW Security VM is the source of the traffic. For traffic to be sent, you must configure a policy rule. Real-time packet flows encapsulated in GRE are sent to the destination IP address that you specify based on the appropriate rule. Traffic matching the policy rule is sent to the destination IP address. Traffic can be sent to either destination IPv4 or IPv6 addresses.

The configuration mirrors the traffic to the external device—it does not imply that traffic is accepted or rejected. You must use subsequent rules in the policy to decide whether to accept or reject the traffic. Mirrored traffic is written also to logs if the rule is configured with **duplicate** in the action field.

You redirect traffic to third-party products for inspection by creating different rules for the type of traffic that you want inspected.

- **Global Settings Rules**—You can configure the vGW Security VM firewall to handle four types of traffic in different ways.
 - IPv6 traffic
 - For homogeneous environments—that is, environments in which all components belong to vGW Series 5.5—this field is grayed out. In this case, you use policy rules to specify whether to allow or deny IPv6 traffic.

- For heterogeneous environments—that is, environments in which one or more components belong to a release prior to vGW Series 5.5—this field allows you to specify whether to allow or drop IPv6 traffic.



NOTE: When you upgrade from vGW Series 5.0 to vGW Series 5.5, the value set for this field is carried over and it determines how IPv6 traffic is to be handled. You can define policy rules to override that behavior.

- Non-IP and Non-ARP traffic

The default firewall configuration drops traffic such as IPX.

- External Logging—vGW Series supports sending logs to third-party Syslog servers. You can enable (or disable) external logging and a Syslog server in this pane. You can also specify from where the logs are sent.



NOTE: The configured external devices must be reachable/routable from the vGW Security VM to every one of the external devices.

By default, the external logging feature is disabled. If you enable it, all traffic that matches a log policy rule, IDS logging, and AntiVirus alerts is written to the vGW Series logs. It is also written to the destination Syslog server that you specify.

You can specify that logging information should be sent to the Syslog server from either of the following locations:

- vGW management server (vGW Security Design VM). Select **Send Syslog from vGW management server**.
- Firewalls (vGW Security VMs). Select **Send Syslog from Firewalls**.

In this case, you can specify whether the firewall logs should also be sent to the vGW management server (vGW Security Design VM). For that purpose, select **Send firewall logs to vGW management server**. By default, they are not sent to the vGW management server.

- NetFlow Configuration—You can specify that all connection flow information is sent through NetFlow Version 9 by enabling the setting in this pane and specifying the IP address and port number of the destination NetFlow collector. Ports 2055 and 9990-9999 are commonly used.



NOTE: Both NetFlow and Syslog are compatible with Juniper Networks STRM.



NOTE: The target server must be reachable/routable from every vGW SecurityVM deployed.

- Infrastructure Configuration Enforcement—VMware requires a special network for communication between the vGW Security VM and VMsafe. This network should not have guest VMs connected to it that are not part of the VMsafe communication process. Enabling this option creates heightened security. If a guest VM that is not part of VMsafe communication is connected to this network and this option is enabled, the VM will be disconnected.

vGW Series IPv6 Support and Global Settings

vGW Series supports both IPv4 and IPv6 addresses. During a transition period, an environment might include components with a mix of both types of addresses.

Externally configured entities, whether a Syslog server, a NetFlow collector, or an external inspection device, *must* be routable over IPv4, IPv6 or both types of infrastructures from every configured vGW Security VM to every configured external entity.

It can happen that some vGW Security VMs might be assigned IPv4 addresses while others might be assigned IPv6 addresses. For example:

- vgw-svm12 might be assigned the IPv4 address 116.27.61.137.
- vgw-svm13 might be assigned the IPv6 address 0810:cdba::3257:9653.

It might be the case that addresses assigned to Syslog, GRE, or Netflow servers and the vGW Security VMs that need to connect to them to send them data belong to different IP protocol families. For example, syslog-server-1 might be assigned the IPv6 address 0680::0202:b3ff:fe1e:8329 whereas vgw-svm12 might be assigned the IPv4 address 116.27.61.137.



NOTE: If a server is identified by its DNS name, the vGW Security VMs will connect and send information to the correct, resolved address: vGW Security VMs with IPv4 addresses will send data to the matching A record and vGW Security VMs with IPv6 addresses will send data to the matching AAAA record.

In cases where the IP protocol families differ, the following results occur:

- vGW Series issues the following message to alert you to the fact that a particular vGW Security VM is not able to connect to the device whose IP address belongs to a different IP protocol family from its own.
- The vGW Security VM does not send out traffic to that server.

For example, if syslog-server-1 is assigned to the IPv6 address 0680::0202:b3ff:fe1e:8329 and vGW Security Design VM vgw-svm12 is assigned to

the IPv4 address 116.27.61.137, vgw-svm1 will not send Syslog messages to syslog-server-1.

- If in the Global Settings page External Logging pane you select the checkbox “Send firewall logs to the vGW management center”, the vGW Security VMs send Syslog messages to the vGW Security Design VM. This behavior applies whether the IP protocol families of the vGW Security VM and the Syslog server match. The Syslog server address has no bearing on sending the logs to the vGW Security Design VM.

However, if you clear this check box at the Security Settings > Security VM Settings page for a particular vGW Security VM, then that vGW Security VM will not send logs to the vGW Security Design VM.

**Related
Documentation**

- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security VM Settings on page 286](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding vGW Series IPv6 Support on page 54](#)

Understanding the vGW Security VM Settings

The Settings module Security Settings > Security VM Settings page allows you to view in one place settings configured for each deployed vGW Security VM. From this page, you can select an individual vGW Security VM and change its configuration, configure a secondary vGW Security VM for it for high availability, and override global settings.



TIP: You can also navigate to this page from the Status section of the Main module. To do so, in the Main module Status page > Status of Security VMs pane, click the row for the vGW Security VM whose configuration you want to view or configure.

The Security VMs pane at the top of the page includes a table with a row for each deployed vGW Security VM, showing the following information:

- Address of the ESX/ESXi host it is deployed to
- Number of VMs that it protects
- If high availability is configured for it
- If network monitoring is enabled
- If NetFlow is configured
- If Syslog is configured
- AntiVirus signature version and data base configuration
- Version of the vGW Security VM

To display a pane that allows you to see detailed information about a vGW Security VM configuration and re-configure its settings, click the vGW Security VM's row.

You can use the tabs shown in the displayed pane to configure unique settings for the vGW Security VM that override the global settings or those set for it when it was installed.

- **VM Settings**—This tab displays configuration information for the vGW Security VM that was set when the vGW Security VM was installed or last modified, in particular the type of IP address assigned to it and how it is obtained. [Figure 137 on page 287](#) shows the pane that you use to change the IP addresses for the vGW Security VM management interface that it uses to communicate with the vGW Security Design VM interface. You can use this pane to override the addresses that were set when you installed the vGW Security VM.

This pane allows you to change the IP protocol family that is used for the vGW Security VM management interface when that protocol does not match that of the vGW Security Design VM with which it must communicate. For information on conditions that would cause an IP address type mismatch between the management interfaces of the vGW Security VM and the vGW Security Design VM, see [“Setting Up vGW Series” on page 36](#) and [“Installing vGW Security VMs on ESX/ESXi Hosts” on page 218](#).

Figure 137: Changing the vGW Security VM Management Interface IP Address

You can change the settings for:

- **IPv4:**

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

- **DHCP**

Use a DHCP server to assign dynamically an IPv4 address to Interface 1. This is the default method. However, this tab page will reflect the configuration set for this vGW Security VM when it was installed or last modified.

- **Static IP**

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

- **IPv6:**

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

- **DHCPv6**

Use a DHCPv6 server to obtain the IPv6 address for Interface 1. This is the default method. However, this tab page will reflect the configuration set for this vGW Security VM when it was installed or last modified.

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

- **Autoconfiguration**

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

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

- **Static IP**

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

- **Network Monitoring**—This tab displays Network traffic monitoring, console monitoring, and NetFlow configuration information.
 - **Network Traffic Monitoring**—By default, network traffic monitoring data is sent to the vGW Security Design VM from all vGW Security VMs. In most cases, it is useful to collect network traffic information that is then displayed in the vGW Series Network module. If you are interested in implementing only firewall protection for VMs protected by this vGW Security VM, you can increase overall system performance by disabling network traffic monitoring on this page. If this option remains enabled for other vGW Security VMs, they will continue to collect and display traffic statistics in the Network module pages.
 - **Console Monitoring**—You can turn console monitoring on or off. Turning it on directs vGW Series to connect to the hypervisor console to monitor traffic in and out of the system to ensure that inappropriate activity is not occurring.
 - **Monitoring off**—No monitoring is performed.
 - **Monitoring on**—Network traffic to the hypervisor console (management center) vNIC is monitored by vGW Series Network module. If netflow is enabled, network traffic information is also available as netflow data.

When Network Traffic Monitoring is disabled, select **Monitoring off** to turn off Console Monitoring. Network Traffic Monitoring must be enabled for Console Monitoring to work.

In addition to monitoring activity as described previously, you can enable IDS traffic monitoring for the vGW Security VM. In this case, network traffic is mirrored to the IDS engine. IDS flags any suspicious activity with high, medium, or low priority alerts, based on how you configured it. For details on IDS, see [“Understanding the vGW Series IDS Module” on page 121](#) and related topics.

To enable IDS monitoring, you must use the IDS tab. If IDS is enabled for the vGW Security VM, the message “IDS inspection console is enabled, see IDS tab.” is displayed in the Console Monitoring box. This message does not appear unless you have enabled IDS.

- **NetFlow Configuration**—You can enable or disable NetFlow for the vGW Security VM. (You cannot change these settings unless NetFlow is enabled globally.)



NOTE: The NetFlow server must be routable from the vGW Security VM.

If NetFlow is enabled, you can direct the vGW Series to send NetFlow data from this vGW Security VM to a different NetFlow collector than the one that is specified in the Settings module's Global section.

To send records to a different NetFlow connector:

1. Select **Enable**.
2. Select **Override global netflow configuration**.
3. Specify the NetFlow collector's address information.
vGW Series supports IPv4 and IPv6 addresses.
4. Click **Save**.

See [Figure 138 on page 290](#).

Figure 138: Configuring Network Monitoring for Individual vGW Security VMs

The screenshot shows the 'Network Monitoring' configuration page. At the top, there are tabs for 'VM Settings', 'Network Monitoring' (which is active), 'IDS', 'Syslog', 'AntiVirus', 'Updates', and 'Support'. A 'Close' button is located in the top right corner. The main content area is divided into three sections:

- Network Traffic Monitoring:** This section has the instruction 'Enable traffic monitoring.' and a checked checkbox labeled 'Enable'. A 'Save' button is at the bottom right of this section.
- Console Monitoring:** This section has the instruction 'Monitor hypervisor console communication.' and two radio buttons: 'Off' (which is selected) and 'Monitor'. A 'Save' button is at the bottom right of this section.
- NetFlow Configuration:** This section has the instruction 'Send records to NetFlow collector.' and two checkboxes: 'Enable' and 'Override global netflow configuration'. Below these is a text input field labeled 'NetFlow collector address:'.

- **IDS**—This tab allows you to enable or disable the IDS engine for the vGW Security VM. If IDS is disabled, you must first enable it globally before you can enable it for the vGW Security VM. You can also turn on IDS inspection of the console by selecting **Enable IDS inspection on the console**.

For details on enabling IDS globally, see [“Configuring Global Settings Using the vGW Series Settings Module \(VMware\)” on page 282](#).

- **Syslog**—This tab allows you to define a Syslog server to use for this vGW Security VM. vGW Series supports sending logs to third-party Syslog servers. To override the Global syslog configuration and specify a different syslog server to use for this vGW Security VM.
 1. Select **Override global syslog configuration**.
 2. Specify the IP address, the port, and the transport protocol of the Syslog server to use for this vGW Security VM.
 3. Click **Save**.
- **AntiVirus**—This tab allows you to enable or disable vGW AntiVirus for this vGW Security VM. If you do not want all ESX/ESXi hosts that are protected by vGW Series to have vGW AntiVirus protection, you can use this page to disable it for the individual vGW Security VM that protects the intended ESX/ESXi host.
- **Updates**—You can use this tab to update the vGW Security VM. For details on updates, see *Understanding the vGW Series Update Settings*.
- **Support**—You can use this tab to enable debug flags to generate debug messages and to collect logs to send to the Juniper Networks Support team for diagnostic purposes. You can also reboot the vGW Security VM from this tab.

Related Documentation

- [Configuring Global Settings Using the vGW Series Settings Module \(VMware\) on page 282](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Understanding and Configuring the vGW Series AntiVirus Settings

This topic explains the vGW AntiVirus settings and how to configure them. Before you read this topic, read "[vGW AntiVirus Configuration Overview](#)" on page 139.

The vGW Security Design VM makes configuration and installation of the vGW AntiVirus feature, including the vGW Endpoint, simple and convenient.

To configure the settings for vGW AntiVirus, you use the AV Settings section of the Settings Module. The AntiVirus Settings page allows you to enable AntiVirus, establish the frequency at which its signature database is updated, and download the vGW Endpoint.

Additionally, a status page displays detailed information, and an **About Juniper vGW Endpoint** box displays the version and build information.



NOTE: When an embedded 30 day license or a license created from the License Management System (LMS) as a Demo license is installed, you can use the vGW AntiVirus feature. However, you cannot update the signatures. That is, the signature updates part of the feature is disabled. In this case, the following message appears beneath the "Current Installed Signatures Version" line: "An appropriate license is required for signature updates".

To update the AntiVirus signatures, a permanent license must be installed.

A vGW Endpoint runs on each protected VM. It is responsible for communicating with the vGW Security VM, monitoring file access, enforcing the AntiVirus policy, and displaying status to the user.

When AntiVirus is disabled, the vGW Security Design VM does not download new signature files, nor will it run On-Demand scans. The vGW Security VM does not load the AntiVirus module, nor does it communicate with the vGW Endpoint.

To enable and configure vGW AntiVirus settings:

1. Check the **AntiVirus Enabled** box.

2. To enable automatic update of the vGW AntiVirus signature database, in the Auto Update section:
 - a. Select **Enabled**.
 - b. Specify the interval in minutes when you want the AntiVirus signature database to be updated automatically.

This section reports the date and version of the currently installed AntiVirus signature database.

To configure the vGW Endpoint and the AntiVirus scan settings:

1. Specify the time after which the vGW Security Design VM should determine that the vGW Endpoint is disconnected.
2. Specify the number of days after which the vGW Security Design VM should consider the current AntiVirus scan outdated.

You can disable vGW AntiVirus from this pane. If you want vGW AntiVirus to remain enabled, but you do not want the AntiVirus signature database to be automatically updated, you can disable automatic updates.

To download the latest version of the vGW Endpoint, click Download. The download section identifies the version and date of the latest vGW Endpoint to allow to you better determine if you want to download it, after you initially download it.

Some administrators download the vGW Endpoint and include it in their boot scripts or software deployment packages that their organization uses. In some cases, organizations place it on a file server. For details on the vGW Endpoint, see [“Understanding and Installing the vGW Endpoint” on page 149](#)

**Related
Documentation**

- [Configuring vGW Series AntiVirus On-Demand Scanning on page 153](#)
- [Configuring vGW Series AntiVirus On-Access Scanning on page 145](#)
- [Understanding vGW Series on page 3](#)

Understanding and Configuring IDS Settings

The Settings module > Security Settings > IDS Settings page allows you to configure IDS settings and IDS updates. See [Figure 139 on page 293](#).

To obtain IDS updates, you must purchase and install an IDS license.

Figure 139: IDS Settings Page

The screenshot shows the Juniper vGW Series Security Settings page. The top navigation bar includes links for Main, Network, Firewall, IDS, Antivirus, Intrusion, Compliance, Reports, and Settings. The left sidebar contains a tree view with categories: vGW Application Settings, Security Settings, and Appliance Settings. Under Security Settings, the 'IDS Settings' option is selected. The main content area is divided into two panels: 'IDS Settings' and 'IDS Updates'.

IDS Settings Panel:

- Intrusion Detection System base settings:**
 - ☒ Enable IDS
- IDS Parameters:**
 - Port numbers for protocols to be treated as HTTP:
 - Port numbers for protocols to be treated as SSL:
- Global Priority Threshold:**
 - Signature priority threshold for enabling signatures by default.
 - ☐ High Priority ☐ High & Medium ☒ All ☐ Custom
- Save** button

IDS Updates Panel:

- IDS Updates:** IDS signatures are updated frequently. The settings below control the behavior of the update processing.
- Update Status:**
 - Currently Installed Signatures: 20130331021143
 - Signatures Available for Update: 20130331021143
 - Last Update Check: Mon Apr 01 12:02:14 PDT 2013
 - Next Update Check: **Check for Update** **Install**
- Automatic Updates (Hourly Check):**
 - ☒ No Automatic Updates
 - ☐ Download Automatically, Manually Apply Updates
 - ☐ Download and Apply Update Automatically
 - Save** button
- Manual Update:**
 - Manually upload an IDS signatures file for processing
 - Browse...** **Clear**
 - Upload File** button

The first time that you set up IDS, you must install the IDS signatures. You use the IDS Updates pane for this purpose. First click **Check for Updates**. After the system searches for the signatures and reports that there are updates, click **Install**.

After the initial signature installation, you can use the Automatic Updates feature or you can manually insert signatures. See [Figure 140 on page 294](#).

Figure 140: IDS Updates Pane

IDS Updates

IDS signatures are updated frequently. The settings below control the behavior of the update processing.

Update Status

Currently Installed Signatures: **20130331021143**
 Signatures Available for Update: **20130428071231**
 Last Update Check: **Wed May 01 20:28:56 PDT 2013**
 Next Update Check:

[Check for Update](#) [Install](#)

Automatic Updates *(Hourly Check)*

☒ No Automatic Updates
☐ Download Automatically, Manually Apply Updates
☐ Download and Apply Update Automatically

[Save](#)

Manual Update

Manually upload an IDS signatures file for processing

[Browse...](#) [Clear](#)

[Upload File](#)

On the IDS Settings page, set the following information to configure IDS for your environment:

- IDS settings
 - Enable IDS—To turn on IDS support, select the **Enable IDS** check box.
 - IDS Parameters— Various ports can be used to pass HTTP and SSL traffic. vGW Series allows you to specify which ports should be analyzed as HTTP and which as SSL.
 - Global Priority Threshold—Specify which signatures are enabled by default based on their priority.
- IDS Updates
 - Update Status—This pane identifies the installed signatures, signatures that are available for updates, and when the last update check was performed. You can check for available signature updates and install them, if any, using this pane as you did for the initial signatures installation.
 - Automatic Updates—Automatic updates are performed hourly. You can enable automatic updates in the following ways:
 - Download Automatically and Manually Apply Updates allows you to apply downloaded signature updates yourself.

- Download and Apply Automatically allows the vGW Security Design VM to apply the updates from vGW Series servers automatically to your local environment.
- Manual Update—You can write or define custom signatures and import them into the vGW Series manually.

**Related
Documentation**

- [Configuring IDS Settings and Viewing Activity on page 128](#)
- [Understanding and Configuring IDS Signatures Settings on page 295](#)

Understanding and Configuring IDS Signatures Settings

This topic explains how to configure settings that control how IDS signatures are managed. You use the IDS Signatures section of the Settings module for this purpose.

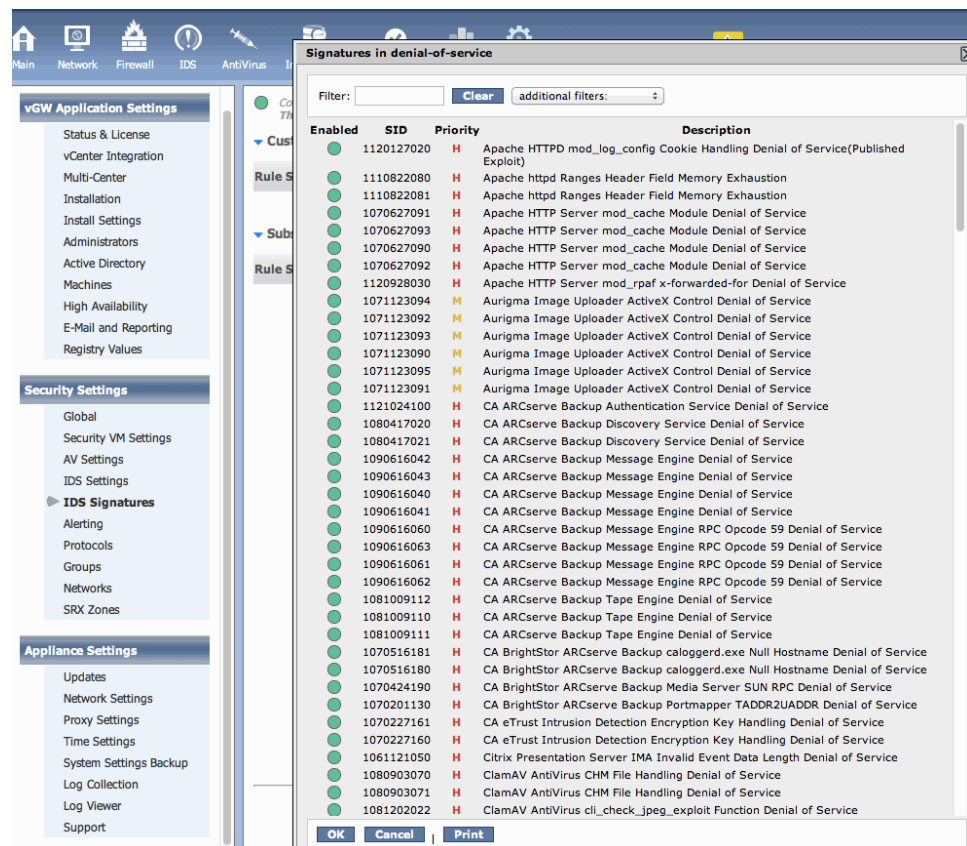
The IDS Signatures page shows the following information:

- **Custom Signatures**, if any have been uploaded. The **Custom Signatures** section does not contain entries until after you manually upload them.
- **Subscription Signatures** shows all signature groups that are part of the standard vGW Series IDS configuration.

You can activate or deactivate entire groups by selecting or deselecting the button under **Rules Selection**.

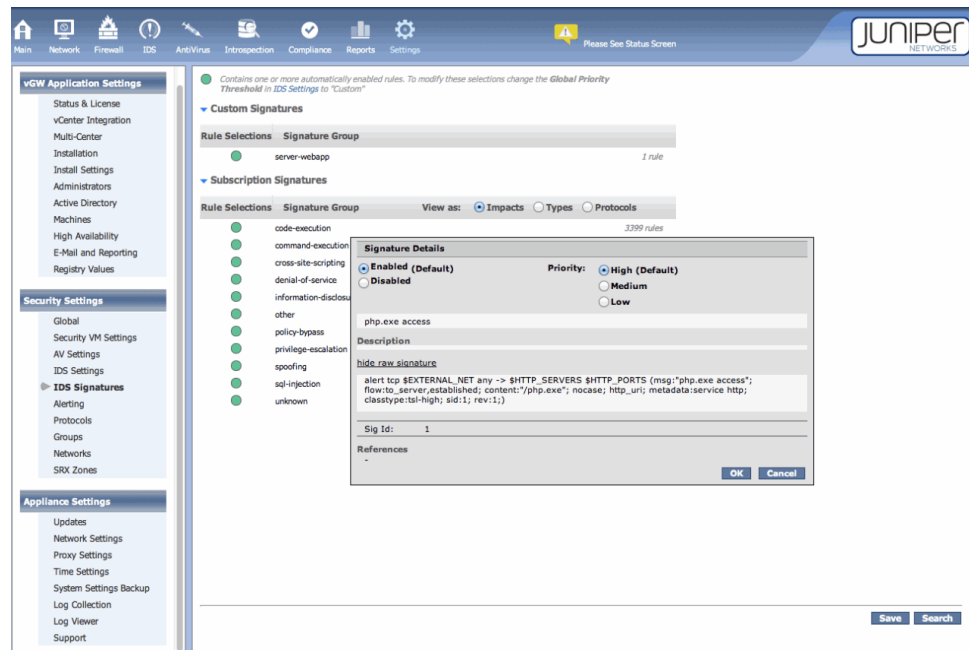
You can view the signatures that belong to a specific group. [Figure 141 on page 296](#) shows the signature rules that comprise the denial-of-service group.

Figure 141: Signatures in a Signature Group



You can also enable or disable individual signatures within a group. To do so, click the signature name in the list of signatures in the group, or the custom signature. Then select the **Enabled** button. The displayed information also provides details on the signature rule. You can change the signature priority level (high, medium, low) on the same Signature Details dialog box. See [Figure 142 on page 297](#).

Figure 142: Signature Details



NOTE: Signature sets that are loaded into the IDS engine apply to traffic that is marked, or tagged, for IDS inspection.

You can manually upload IDS custom signatures using the IDS Signatures section of the vGW Security Design VM Settings module.

Related Documentation

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

Understanding the vGW Series Security Alert Settings

This topic covers the events sent by e-mail and SNMP and how to configure Alert settings for e-mail and SNMP traps.

It includes the following sections:

- [Event Types on page 297](#)
- [E-mail Alert Settings on page 298](#)
- [SNMP Trap Settings on page 298](#)
- [AutoConfig and Multicast Alerts on page 298](#)

Event Types

vGW Series sends security alerts (Main→Events and Alerts→Security Alerts) by e-mail and SNMP. Security Alerts have high, medium, and low (H/M/L) priorities. By default, alerts of all priorities are sent by SNMP and e-mail. However, you can use the center.conf

parameter `center.alert.notification.priority` to change this configuration. By default, it is set to 3 (low). Alerts with a priority that is equal to or lower than the configured value are sent.

E-mail Alert Settings

You enable e-mail alerts by providing the mail relay server IP address and the source and destination e-mail addresses. vGW Series supports both IPv4 and IPv6 addresses. The aggregation time is the gap between successive notifications.

You do not need to configure multiple e-mail recipients. However, you can create four custom e-mail alert tags that point to different e-mail aliases or individual e-mail accounts, or a combination of the two. You can specify these custom tags in the security policy editor.

To send both an e-mail alert and an SNMP trap on a single rule, you use the standard alert icon. In this case, only the e-mail addresses listed in the **Recipients Addresses** are used. That is, you can not use custom tags when you send e-mail and SNMP alerts.

SNMP Trap Settings

An Simple Network Management Protocol (SNMP) trap is an asynchronous notification from agent to manager. It includes the current `sysUpTime`, and OID identifying the type of trap, and optional variable bindings. SNMP traps can be set via Version 1 or Version 2. You must enter the SNMP server address and community string. Optionally, you can set the aggregation time again (the delay between successive events).

To configure SNMP using the Settings module Alerting > SNMP Trap Settings pane:

- Select the **Enable** check box if you want to send SNMP traps on alerts.
- Select the SNMP version to use. By default Version 1 (SNMPv1) is selected.
- Specify the SNMP monitor address.
- Specify the SNMP community string.
- Specify the aggregation time in seconds.
- Click **Save**.

AutoConfig and Multicast Alerts

By default, the vGW Series is configured to alert when autoconfig addresses are discovered (Settings page -> 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 it can not acquire a DHCP lease, it defaults to using an autoconfig address in the 169.254.*.* range. Often this setting represents a configuration problem or an issue with the DHCP service.
- **Multicast**—Many hosts use multicast packets to advertise their presence on the network. They also send broadcast information about the services that 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](#)

Understanding vGW Series Protocols Support

This topic contains the following sections:

- [The Protocols Page and Table on page 299](#)
- [Creating Protocol Groups on page 299](#)
- [ICMPv6 on page 299](#)
- [Additional Protocols Added for IPv6 on page 300](#)

The Protocols Page and Table

The Settings module Security Settings > Protocols page lets you view existing available protocols and add to the list. By default, the protocols table shows all IANA registered protocols. You can add to this table custom protocols or other application protocols that are not IANA registered. Protocols that you add are shown by name in network reports instead of being displayed by port or protocol.

You can also define your own non-TCP, non-UDP and non-ICMPv6 protocols such as GRE and IPsec protocols. You can define protocol ranges such as Custom App /TCP/8000-8005.

The Protocols table displays the name of a protocol, the kind of protocol it is, such as TCP, UDP, or ICMPv6, and the number of the port used for it or the type if it is an ICMPv6 protocol.

Creating Protocol Groups

You can combine a number of protocols into a Protocol Group so it can be used in Firewall Policy creation (for example, policies for Global, Group, or Individual VMs can include it).

To do so, click **Add**, enter a name for the group, select the appropriate protocols, and click **Save**.

ICMPv6

The protocols table includes the following Internet Control Message Protocol version 6 (ICMPv6) protocols:

- 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.

ICMPv6 is integral to IPv6 and fundamental to the proper functioning of IPv6 networks. For details on how the vGW Series firewall handles ICMPv6 protocols and the default protocol group for ICMPv6 protocols, see [“Understanding How vGW Series Handles ICMPv6 Protocol Traffic” on page 106](#).

Additionally, vGW Series ICMPv6 (icmpv6) was added as a new transport protocol type. To use it, you must specify in the **Type:** box the protocol type number.

Additional Protocols Added for IPv6

You can use the following IPv6 protocols, which were added to the supported protocols in firewall policies:

- dhcp6-client
- dhcp6-server
- ripng
- route-ipv6
- frag-ipv6
- nonxt-ipv6
- opts-ipv6

Related Documentation

- [Understanding vGW Series IPv6 Support on page 54](#)
- [Understanding the vGW Series Firewall Module on page 93](#)
- [Configuring vGW Series Firewall Policies on page 114](#)
- [Understanding How vGW Series Handles ICMPv6 Protocol Traffic on page 106](#)
- [Understanding vGW Series on page 3](#)

Understanding vGW Series Groups

The Settings module **Security Settings > Groups** page lets you define groups that can contain VMs and resources. You can automate many security tasks by putting in place the proper group structure.

This topic includes the following sections:

- [Uses of Groups on page 301](#)
- [vGW Series Group Types on page 301](#)
- [Policy Groups and Monitoring Groups on page 301](#)
- [Defining the Group as a Policy Group Option with Automatic or Manual Selected on page 302](#)
- [Copying Groups on page 302](#)

Uses of Groups

Groups serve many purposes. For example, you might want to use a group for the following reasons:

- To understand how VMs belonging to the group interact on the network, but you do not want to protect their traffic.
- To use the group as source or destination term of a firewall policy rule.
- To apply policy to the VMs of a group automatically. In this case, when you create the group, you define it as a Policy Group.
- To check the compliance of VMs that belong to a group.

vGW Series Group Types

vGW Series supports the following two group types:

- *Static Groups* that allow you to define a collection of objects, such as networks, VMs, or external physical systems. A static group remains the same unless you manually change it.
- *Smart Groups* that allow for the dynamic association of VMs. To create a Smart Group, you define a set of rules, or requirements, that specify variables that characterize the group. A VM that matches one or more of a Smart Group's variables automatically becomes a member of the group.

A VM may pass in and out of Smart Groups automatically as the VM's configuration changes. Without your interaction, a VM that matches one or more of a group's variables, based on its rule requirements, is inserted into the Smart Group. If the VM's configuration is changed in such a way that it no longer meets the group's definition, the VM is automatically removed from it.

When the VM enters the group, the group's policy is applied to it. When it leaves the group, the group's policy is removed from it.

Policy Groups and Monitoring Groups

You can select the Policy Group option when you define a group to control policy association. When you select the Policy Group option, the group shows up in the Policy Groups area of the VM tree.

Groups that do not have a policy associated with them appear by default in the Monitoring Groups section of the VM tree.

The VM tree contains:

- **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 without the Policy Group option selected, 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.

For details on how to enable Firewall Monitoring, see [“vGW Series VMsafe Firewall + Monitoring and VMsafe Monitoring Modes” on page 22](#).

- Monitored/Secured VMs—Lists VMs monitored by the vGW Series, VMs that have a firewall protecting their network traffic, or both.

Defining the Group as a Policy Group Option with Automatic or Manual Selected

You can select the Policy Group option when you define a group to control policy association. When you select the Policy Group option, the group shows up in the Policy Groups area of the VM tree. Groups that do not have a policy associated with them appear in the Monitoring Groups section of the VM tree.

To define a policy for the group, you use the Firewall module, select the group in the VM tree, and configure its policy rules. To install the policy, you use the Firewall Module > Apply Policy page.

Among the information that you configure for a group that you define as a Policy Group is how the policy is applied:

- Automatic—Policy changes for the group's VM members occur without your intervention. That is, you do not need to use the Firewall module's Apply Policy page to push the policy out to the VMs.
- Manual—You manually apply a policy to the VMs that belong to the group.

When you add a VM to a Smart Group or a Static Group, or a VM matches a Smart Group attribute and enters the group because of the match, the VM gets the policy rules associated with the group if the following conditions are met:

- The group is configured as a Policy Group and there is a policy containing policy rules associated with it.
- The group is configured with the Automatic option selected.

[“Automatically Applying Policy Rules to VMs in Policy Groups” on page 303](#) gives details on defining Smart Groups and Static Groups as Policy Groups with the Automatic Option to automatically “push” policies to VM members of a group.

Although a VM that enters a group—either because you added it or dynamically because it matched a Smart Group variable—gets the group's policy, this will not start to occur until after the first use.

Also, if changes are made with the vGW Series Cloud SDK, you must apply them either using the vGW Security Design VM Firewall module > Apply Policy page or using the relevant function. They do not take effect simply because the vGW Security Design VM is changed.

Copying Groups

You can use the Group page to duplicate groups.

To copy groups:

1. From the Settings module on the vGW Security Design VM, select the **Security Settings > Global**.
2. In the Groups table, click the name of the group that you want to copy.
3. Click **Copy Group**. A dialog box appears.
4. Give the new group a name.
5. If the group that you are copying is a policy group, click **Keep Policy** if you want the original group's policy to be associated with the new group.
6. For a Smart Group, you can:
 - Click **Duplicate Smart Group logic** to duplicate the rule set on the copy.
 - Click **Convert VM membership to static group** to create a static group that contains the members of the copied Smart Group.
7. Click **Save**.

The new group is added to the Groups table.



NOTE: A new group created as a copy inherits the auto push property of the original. However, because it is effectively a new group, it must be manually pushed initially.

Related Documentation

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

Automatically Applying Policy Rules to VMs in Policy Groups

vGW Series allows you to create Static Groups or Smart Groups that are defined as Policy Groups and then associate policy rules with them. If you select the Automatic option for the group when you configure it, when a VM joins the group, the policy rules associated with the group are automatically applied to the VM.

Configuring a group as a Policy Group whose rules are applied automatically to its VM members entails:

- Creating either a Static Group or a Smart Group and selecting the Policy Group and Automatic options. Use the Settings > Security Settings > Groups page.

The Groups page allows you to define a Static Group or a Smart Group and attributes for it. You can specify that the group is a Policy Group and you can select Automatic for it. If you select Automatic, rules defined for the policy group are applied automatically to VMs that join the group. You can add VMs to a policy group (Static Group) or they join it dynamically because they match one or more of the group's configured variables (Smart Group). In either case, when the VM joins the group, it gets the group policy rules.

- Securing the group with a firewall policy. Use the Firewall module to create the policy rules for the group and apply (install) the policy rules.

When you add a VM to a Smart Group or a Static Group or a VM matches a Smart Group attribute and enters the group because of the match, the VM gets the policy rules associated with the group if the following conditions are met:

- The group is configured as a Policy Group and there is a policy containing policy rules associated with it.
- The group is configured with the Automatic option selected.

After you create a policy group, it is added to the Firewall > Apply Policy table, ready to be applied to the group. The policy group must be applied once before it can be used for auto-push. Note that auto push will apply the policy to a VM automatically only when the VM enters or exits the group based on matching.

This example shows how to create a group that automatically pushes its policy to VMs that belong to the group or join it dynamically. It creates a Smart Group called HighPriorResGrp, and it configures it as a Policy Group with the Automatic option selected.

1. On the Settings > Security Settings > Groups page, configure a Smart group called HighPriorResGrp that watches for any VMs connected to a particular VMware resource pool (called high-prior-res) obtained through vi.resourcepool.

Smart Groups specify attributes that a VM must match to join the group. For details on Smart Groups, see [“Understanding vGW Series Smart Groups” on page 305](#).

When a VM joins the Policy Group, the group's rules are instantly installed on that VM without requiring any intervention on your part. [Figure 143 on page 304](#) shows the Smart Group configuration. Notice that **Policy Group** and **Automatic** are selected for the Group Attributes.

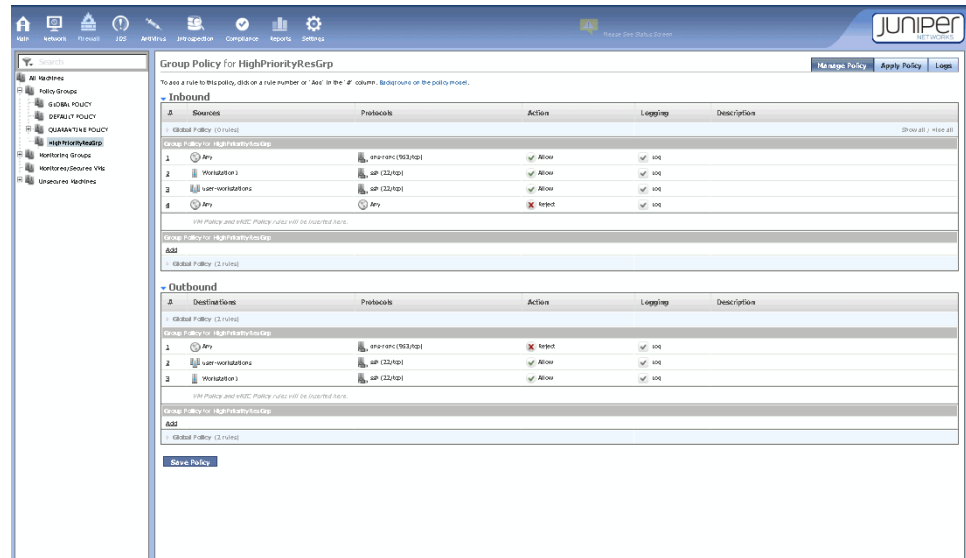
Figure 143: Configuring a Smart Group As a Policy Group

2. Configure policy rules for the HighPriorityResGrp Smart Group.

When you create a group and define it as a Policy Group, vGW Series places it in under Policy Groups in the VM Tree. You can click on the group name to display the Firewall

> Manage Policy tab that allows you to configure group rules. See [Figure 144 on page 305](#).

Figure 144: Configuring Policy Rules for a Smart Group with Policy Group Enabled



Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security VM on page 68](#)
- [Understanding the vGW Series Application Settings on page 206](#)
- [Understanding Policy per vNIC and Smart Groups for VMware Environments on page 266](#)

Understanding vGW Series Smart Groups

This topic includes the following sections:

- [Background on page 305](#)
- [About Smart Groups on page 306](#)

Background

In most organizations the virtualized environment changes rapidly. Environmental change cycles that occur over days and weeks in the physical data center take place almost instantaneously in the virtualized environment. New virtual machines (VMs) can be cloned or created from templates; new virtual hardware can be added to existing VMs and reconfigured; and existing VMs can be moved from one network to another in a matter of minutes or even seconds.

To accommodate these rapid changes and to allow you to secure the network during them, the vGW Series provides a feature called Smart Groups. A Smart Group is characterized by a set of rules that set the membership criteria. If a VM's configuration

matches those rules, it is dynamically associated with the Smart Group. Furthermore, if you define the Smart Group as a policy group, the policy is automatically pushed out to the VMs in the group and those that are added to it.

Smart Groups allow you to maintain complete control. Security changes can be applied automatically and instantly, or simple alerts can be generated signaling the need for manual intervention.

With Smart Groups:

- Security policies can adapt instantly to changes in the virtualized data center without manual intervention.
- Compliancy checks ensure that security is maintained and that risks are mitigated immediately.

About Smart Groups

To create a Smart Group, you define one or more rules created as expressions. For a VM to belong to a group, its configuration must meet one or more of the rules' criteria, depending on your specifications. You can specify whether a VM is required to meet all the criteria or only part of it.

Smart Groups are dynamic in that their membership can change rapidly. VMs can be added to or removed from Smart Groups automatically within seconds. At any time a VM's configuration might be changed in a way that now causes it to match a Smart Group. If a VM no longer matches a Smart Group's rules, it is removed from the group. You can observe this transition in the VM tree. A VM within a Smart Group appears within the group under Policy Groups. When that VM no longer matches the Smart Group's rules, it is moved to Monitoring Groups.

Related Documentation

- [Understanding vGW Series Groups on page 300](#)
- [Understanding vGW Series on page 3](#)
- [About Using vGW Series Attributes for VMware on page 306](#)

About Using vGW Series Attributes for VMware

vGW Series continuously analyzes both its own and the VMware objects databases in relation to the Smart Group rules that you configure to determine if a VM should belong to a Smart Group or not. The rules that you configure to define a Smart Group are obtained from these locations:

- vGW Series Smart Group attributes. These attributes are categorized and labeled with the prefix *vf*.
- VMware vCenter attributes. These attributes are labeled with the prefix *vi*. VMware's vCenter identifies attributes such as the port group to which the virtual network interface is connected.

- VMware vCloud Director attributes. These attributes are labeled with the prefix *vcd*. The vGW Security Design VM obtains metadata associated with a VM from the vCloud Metadata tab page for the VM.

You can associate Smart Groups with a firewall policy. Policy association is controlled by the Policy Group option that you can select when you define the Smart Group.

Using Smart Groups, you can streamline policy application to ensure security efficiently throughout your virtual infrastructure. Firewall policies are applied to VMs instantly without your intervention when a VM becomes a member of a Smart Group. Consider these two cases in which firewall policies are automatically applied to VMs:

- Suppose that you associate the virtual network interface of a VM with the corporate production network. As a consequence of the configuration change, the VM meets a Smart Group's rule that specifies the *vi.portgroup* attribute and matches the configuration. In this case the VM becomes a member of the Smart Group, and the Smart Group's firewall policies are applied to it.
- Suppose that you define a Smart Group that checks for VMs connected to a particular VMware resource pool that is specified in a rule that uses the *vi.resourcepool* attribute. When you add a VM to this resource pool, the VM is added to the Smart Group and the Smart Group's firewall policies are applied to it.

Related Documentation

- [Understanding vGW Series Groups on page 300](#)
- [Understanding vGW Series on page 3](#)
- [Creating vGW Series Smart Groups for VMware on page 307](#)
- [vGW Series Attributes for VMware on page 311](#)

Creating vGW Series Smart Groups for VMware

This topic explains how to configure vGW Series Smart Groups. You can create groups comprised of members who meet or violate the designated match criteria defined in the Matches field of the Smart Group.

To define a Smart Group, you use the Settings module Security Settings > Groups page, and click **Add Smart Group**. The editor has two modes: Basic and Advanced. The default mode is Basic.

Suppose you want to create a compliance rule that states that all Web server VMs should have version Apache 2.x installed because of known security issues in versions 1.x. You can configure a Smart Group for a compliance rule and configure vGW Series to issue an alert when any Web server currently in production or brought online in the future has a version of Apache that is prior to 2.x.

Smart Group creation options—the parameters used to define the group—are obtained from two locations: namely, Security Design vGW attributes and vCenter attributes. Through VM Introspection, the vGW Security Design VM can discover items such as which applications are installed on a VM, while VMware's vCenter identifies attributes such as the port group to which the virtual network interface is connected. There are numerous

attributes each classified into “vf” (vGW-based) and “vi” (vCenter-based) categories as described in the topic [“vGW Series Attributes for VMware” on page 311](#).

The following values are returned for the Type field.

- Boolean: True or False
- Integer: Numeric value
- String: Free-form text string
- Multi String: Multiple string values concatenated together with separators such as commas, semicolons, or slashes
- Multi Value: List of available choices

In Basic mode you can select one or more attributes and assign an **All** or **Any** constraint. You add rules by clicking the + sign.

[Figure 145 on page 308](#) shows a group called WebServers that is created when the VMware vCenter name (vi.name) contains www and the application named Apache is installed on the VM. Both conditions must exist for a VM to be included in this group. The information that defines this Smart Group is obtained through VI Introspection and is stored in vf.application.

Figure 145: Creating a Smart Group Using Basic Mode

Add Group

Name:

Advanced

Matches: ☒ All ☐ Any

vi.name	Contains	www	-	+
vf.application	Contains	apache	-	+

Group Policy Attributes:

☒ Policy Group

Priority:

Apply Policy: ☐ Automatic ☒ Manual

Test Save Cancel

To define a Smart Group using basic mode:

1. Select **Setting > Security Settings > Groups**.
2. To create a new Smart Group:
 - a. Click **Add Group**.
 - b. On the displayed pane, click **Add a Smart Group**.

If you do not know the meaning of an attribute or the values that it can take, click ? at the end of the row. The pop-up message box that appears describes the attribute. It gives its data type, and it identifies possible values.

3. Give the Smart Group a short, descriptive name. The name is displayed in the Groups table.
4. For Matches, select **All** if the VM must meet all criteria defined in the field below or **Any** if the VM can meet any of the criteria defined in the field below.
5. For each row, select the following information:
 - An attribute.
 - A comparator. For example, you can require that a VM must meet the attribute specification to be associated with the group, or you can define a rule that excludes VMs that meet the criteria.
 - A value.
6. Select the **Policy Group** check box if you want the Smart Group to belong to a policy group.

When you select Policy Group:

- The Smart Group is added to the Policy Groups area in the VM tree.

You can now configure a firewall policy for the Smart Group on its Group Policy page. You use the Firewall module in conjunction with the VM tree to display the Smart Group's policy page.

- Specify a priority level and a precedence level:
 - You can select high, medium (default), or low for the priority level.
 - You can use **Precedence** within **Level** to define the precedence for Smart Groups that are created with the same priority level.



NOTE: A VM can belong to more than one Smart Group. In this case, the policy rules of all Smart Groups that the VM is a member of are applied to the VM. How the rules are applied also depends on the precedence and priority settings.

It can happen that more than one Smart Group is defined with the same priority level and the same precedence within that level. In this case, Smart Group rules are applied to the VM in the order in which the Smart Groups were created.

7. Test the configuration before you save the Smart Group definition. Click **Test** to verify that the group contains the VMs that you intended it to include.

In addition to creating a Smart Group by adding rows to the rules table using Basic mode, the editor's Advanced mode allows you to write regular expressions to construct more

complicated scenarios. [Figure 14-6 on page 310](#) shows how to define the simple WebServers example in Advanced mode using a regular expression.

Figure 14-6: The Smart Group Editor in Advanced Mode Using Regular Expressions

Add Group

Name:

Basic

Selection query:

Group Policy Attributes:

☒ Policy Group

Priority:

Apply Policy: ☐ Automatic ☒ Manual

The selection query allows you to define expressions based on a simple set of operators. You can write an expression in the context of each VM, getting its attributes, and if the expression evaluates as True, the VM becomes part of the group. [Table 18 on page 310](#) covers the various Smart Group attribute types and operators.

Table 18: Operators for Creating Smart Groups Using Regular Expression

Attribute Type	Supported Operators
String	<p>The most common attribute type.</p> <p>Contains (~), Not-Contains (!~), Equals (=), Not-Equals (!=), Matches RegExp (=~).</p> <p>Full wildcard support such as name = "finance-*" is recognized.</p>
Numerical	Equals (=), Greater than (>), Not-Equals (!=), Less-Than (<), In (in), Not in (not_in).
IP	Equals (=), In (in), Not in (not_in).
Boolean	Equals (=), Not-Equals (!=) Return value is either true or false. For example, vf.secured = false or vf.secured != true.
Multi	Contains (~), Not-Contains (!~), Equals (=), Not-Equals (!=), Matches RegExp (=~).
Group	Contains (~), Not-Contains (!~), Equals (=), Not-Equals (!=), Matches RegExp (=~).

You can also create wildcard matches if you match on a full string. For example:

- *.WWW.* - Match VMs with WWW anywhere in the name"
- ^Corp.* - Matches VMs starting with Corp

- `.*1$` - Matches VMs ending with "1".
- `.*Tier-[1-3].*` - Match VMs with Tier-1/2/3"
- `^[ABC].*` - Match VMs starting with A, B, or C.

**Related
Documentation**

- [About Using vGW Series Attributes for VMware on page 306](#)
- [vGW Series Attributes for VMware on page 311](#)
- [Understanding vGW Series Groups on page 300](#)
- [Understanding vGW Series on page 3](#)

vGW Series Attributes for VMware

Table 19 on page 311 identifies the attributes that you can use in defining Smart Groups.

Table 19: Smart Group Attributes

Attribute name	Data Type	Description
<code>vcd.tag</code>	String	vCloud Director Organization and metadata attributes.
<code>vf.antivirus.database.version</code>	String Value	What version of AV database version is this VM using? (What's installed on the central AV database it is connected to)?
<code>vf.antivirus.endpoint.connected</code>	Boolean Value	Is this VM properly connected to central AV scan engine?
<code>vf.antivirus.endpoint.enabled</code>	Boolean Value	Does this VM have an operational AV agent installed?
<code>vf.antivirus.endpoint.version</code>	String Value	Version of endpoint installed on the VM.
<code>vf.antivirus.engine.version</code>	String Value	What version of the AV engine is this VM is using? (What is installed on the central VM database it is connected to?)
<code>vf.antivirus.onaccess.enabled</code>	Boolean Value	Does this VM have on-access AV scanning enabled?
<code>vf.antivirus.quarantine.enabled</code>	Boolean Value	Is this VM configured to quarantine virus files?
<code>vf.app_count_bad</code>	Integer	Number of applications on a VM that are classified as bad.
<code>vf.app_count_known</code>	Integer	Number of applications on a VM that are classified as known.
<code>vf.app_count_unclassified</code>	Integer	Number of applications on a VM that are unclassified.

Table 19: Smart Group Attributes (*continued*)

Attribute name	Data Type	Description
vf.app_count_unknown	Integer	Number of applications on a VM that are classified as unknown.
vf.app.gi.compliant	String Value	Is this VM in compliance with the selected Gold Image?
vf.app.is.gold.image	Boolean Value	Is this VM defined as a master image for Image Enforcer comparisons?
vf.app.matches.gold.image	Boolean Value	Is this VM compliant with its configured Gold Image?
vf.app.registry	String Value	Registry value from s registry as determined by Introspection of VM.
vf.application	String Value	An application installed on a VM.
vf.description	String	The text string description of the VM, as defined in the vGW Security Design Settings module Machines section.
vf.firewall	String	Is this VM a vGW Security VM?
vf.group	Multi String	Comma-separated string of all vGW groups to which a VM belongs.
vf.has_installed_group_policy	Boolean	Does the VM have a non-default group policy installed?
vf.has_installed_policy	Boolean	Does the VM have an installed security policy?
vf.hotfix	Multi String	Hotfix installed on a VM.
vf.monitored	Boolean	Is the VM currently being monitored by the vGW Security Design VM?
vf.name	String	Name as defined in the vGW Security Design VM.
vf.os	String	The operating system installed on the VM.
vf.quarantined	Boolean Value	Is this VM in a quarantined state, and thus in the Quarantine Policy group?
vf.secured	Boolean	Is a VM currently secured by the vGW Security Design VM?
vf.secured_active	Boolean	Is the VM actively protected by vGW?

Table 19: Smart Group Attributes (*continued*)

Attribute name	Data Type	Description
vf.tag	String	Tags associated with this VM that are semicolon separated.
vf.type	Enumeration	The machine object type.
vf.virus.infected	Boolean Value	Has a virus been detected on this VM by the vGW antivirus engine?
vi.attribute	String Value	The attribute values that are defined in the annotation box in VI.
vi.cluster	String	Cluster containing a VM.
vi.datacenter	String	Data Center in vCenter where a VM is housed.
vi.deleted	Boolean Value	Has this VM been deleted?
vi.excfg.copy.disable	Boolean Value	Is the copy and paste to remote console feature disabled for this VM?
vi.excfg.deviceconnectable.disable	Boolean Value	Is this VM configured to allow devices to be connected?
vi.excfg.deviceedit.disable	Boolean Value	Is this VM configured to allow devices to be connected and removed?
vi.excfg.diskshrink.disable	Boolean Value	Is this VM configured to prevent virtual disk shrinking?
vi.excfg.diskwiper.disable	Boolean Value	Is this VM configured to prevent virtual disk shrinking?
vi.excfg.dragndrop.disable	Boolean Value	Is the copy and paste to remote console feature disabled for this VM?
vi.excfg.hostinfo.disable	Boolean Value	Is access to host performance information available to this VM?
vi.excfg.log.disable	Boolean Value	Is the VM log file size limited for this VM?
vi.excfg.log.keep.old	Numeric Value	Is the number of stored log files limited for this VM?
vi.excfg.log.rotatesize	Numeric Value	Is the VM log file size limited for this VM?
vi.excfg.paste.disable	Boolean Value	Is the copy and paste to remote console feature disabled for this VM?

Table 19: Smart Group Attributes (*continued*)

Attribute name	Data Type	Description
vi.excfg.remotedisplay.max	Numeric Value	How many remote consoles are available for this VM? VMware Hardening guideline recommends limiting to one.
vi.excfg.remoteop.disable	Boolean Value	Are remote operations disabled for this guest?
vi.excfg.setguiopts.disable	Boolean Value	Is the copy and paste to remote console feature disabled for this VM?
vi.excfg.vmxfilesize.limit	Numeric Value	Is the VMX file size limited (to limit the informational messages from VM to VMX file) ?
vi.folder	Multi-String	The folder containing a VM in vCenter.
vi.host	String	ESX/ESXi hosting a VM.
vi.host.console.ids	Boolean Value	Is vGW IDS inspection enabled for this hypervisor's service console?
vi.host.console.monitor	Boolean Value	Is vGW network monitoring enabled for this hypervisor's service console?
vi.host.lockdown	Boolean Value	Is lockdown mode enabled for this hypervisor host?
vi.host.ntp.enabled	Boolean Value	Is Network Time Protocol (NTP) configured and enabled for this hypervisor?
vi.host.techsupportmode.disable	Boolean Value	Is tech support mode enabled for this hypervisor?
vi.host.vmkernel.isolated.vlan	Boolean Value	Is the vmkernel management network on this hypervisor on an isolated VLAN?
vi.host.vmkernel.isolated.vswitch	Boolean Value	Is the vmkernel management network on this hypervisor on an isolated vSwitch?
vi.indep.nonpersist.disk.ct	Numeric Value	The number of virtual disks used by this VM that are configured as Independent nonpersistent and thus cannot be introspection scanned.
vi.ipv4	IPv4 (multi value)	The IP addresses as known on a VM.

Table 19: Smart Group Attributes (*continued*)

Attribute name	Data Type	Description
vi.ipv6	IPv6 (multi value)	<p>The IP addresses as known on a VM. They can be coded as single addresses or an address range.</p> <p>Example Addresses:</p> <ul style="list-style-type: none"> • 2001:0db8:85a3:0000:0000:8a2e:0370:7334 • fe80::202:b3ff:fe1e:8329
vi.memory_inspection	Boolean	Are VMsafe memory and CPU API enabled for this VM?
vi.name	String	Name of this VM as defined in vCenter.
vi.notes	String	Annotation free text notes attached to the VM in vCenter.
vi.os	String Value	Operating system defined for the VM in vCenter.
vi.pg.security.forgedtransmits	Boolean Value	Is VM connected to a port group that allows forged MAC addresses (MACs other than defined in the VMX)?
vi.pg.security.macchanges	Boolean Value	Is VM connected to a port group that allows reception of unknown MAC addresses (MACs other than defined in the VMX)?
vi.pg.security.promiscuous	Boolean Value	Is VM connected to a promiscuous port group?
vi.portgroup	String Value	Port groups on the virtual switch this VM is actively connected to. Port Groups for disconnected vNICs will not be included. (For a running/suspended VM, this will be the port groups actually connected. For a stopped VM, this value is the port groups that are connected at poweron.)
vi.portgroup.all	String Value	Port groups on the virtual switch this VM is connected to. This list includes port groups even if the vNIC is disconnected. (For a running/suspended VM, this will be the port groups actually connected. For a stopped VM, this value is the port groups that are connected at poweron.)
vi.powerstate	Enumeration	What is the current power state of this VM?
vi.pvlan	Numeric Value	Private VLAN values for connected port groups.
vi.pvlan.all	Numeric Value	List of all private VLANs in use by this VM, includes vNICs in both connected and disconnected states.

Table 19: Smart Group Attributes (*continued*)

Attribute name	Data Type	Description
vi.os	String	Operating system defined for the VM in vCenter
vi.resourcepool	String	Resource pool VM is a member of vCenter.
vi.snapshots.count	Numeric Value	How many snapshots exist for this VM?
vi.vapp	Multi String	vApp group VM is a member of vCenter.
vi.vlan	Multi-value integer	VLANs of connected port groups.
vi.vlan.all	Multi-value integer	VLANs of all interfaces.
vi.vmci_enabled	Boolean	Is VMCI (shared memory communications) enabled for this VM?
vi.vmsafe_configured	Boolean	Is VMSafe firewall security enabled for this VM?
vi.vmsafe_dvfilter	Multi String	The dvfilters protecting this VM.
vi.vmsafe.initfailmode	Enumeration	If VMSafe is unable to initialize, what is the network connectivity choice for this VM?
vi.vmwaretools.running	Boolean	Is VMware Tools running on this VM?
vi.vmwaretools.uptodate	Boolean	Is the version of VMware Tools installed on this VM current?
vi.vnic.count	Numeric Value	Number of connected vNICs.
vi.vswitch	Multi String	vSwitch VM is connected to.

Related Documentation

Understanding the Settings Module Networks Settings

You can use the Settings module **Security Settings > Networks** page to define network objects for use in vGW Series firewall policies. You can define a network by IP Range or Subnet Mask. vGW Series supports IPv4 and IPv6 environments, and, as such, network ranges in both address spaces.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series IPv6 Support on page 54](#)
- [Configuring the vGW Series Policy per vNIC Feature on page 261](#)

Understanding the vGW Series SRX Zones Settings

You can use the SRX Zones section of the Settings module of the vGW Security Design VM to create interoperability with physical SRX Series devices. For details, see [“vGW Series and SRX Series Security Zones” on page 359](#). vGW Series integration with SRX Series zones allows it to obtain zone information from an SRX Series device and populate the vGW Security Design VM with that information. vGW Series can also put VM information into SRX Series address books that allows you to know which VMs are mapped to each zone.

Related Documentation

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

vGW Security Design Appliance Settings

- [Updating the vGW Security Design VM on page 319](#)
- [Updating Individual vGW Security VMs on page 321](#)
- [Updating vGW Security VMs in Batch Mode on page 322](#)
- [Configuring the vGW Series Network Settings on page 323](#)
- [Configuring vGW Series Proxy Settings on page 328](#)
- [Configuring vGW Series Time Settings on page 328](#)
- [Understanding the vGW Series Backup and Restore Feature on page 329](#)
- [Configuring the vGW Series Backup and Restore Feature on page 331](#)
- [Understanding vGW Series Log Collection on page 333](#)
- [Viewing the vGW Series Logs on page 336](#)
- [Understanding vGW Series Support Settings on page 336](#)

Updating the vGW Security Design VM

This topic covers how to update the vGW Security Design VM online and offline manually.

This procedure explains how to update the vGW Security Design VM online. In this case, the vGW Security Design VM must be able to connect to the Juniper Networks update servers (HTTPS - TCP 443).

To update the vGW Security Design VM manually:

1. Ensure that a proper entitlement key is installed on the vGW Security Design VM.

If it is not already installed, insert the entitlement key in the vGW Application Settings -> Status & License section of the Settings module. This section also allows you to see the update status of the vGW Security Design VM

- a. In the Product Licensing section, click **Manage Licenses**.

A table showing the installed licenses is displayed.

This section also allows you to see the update status of the vGW Security Design VM.

Without an entitlement key, you cannot activate and install an update. You obtain the entitlement key when you purchase the product and software subscription contract.

2. Navigate to Settings > Appliance Settings > Updates.

- a. Click **Check for Updates** to query the Juniper Networks update servers.

The update server checks to determine if the component requires an update.

- b. If an update exists, click **Update Now** to apply the changes.

vGW Series downloads the required updates from the Juniper Networks server. In some cases, the vGW Security Design VM will need to be rebooted.

Updating the vGW Security Design VM Offline

This procedure explains how to update the vGW Security Design VM offline, that is, without using Internet access. It uses an ISO image connected to the vGW Security Design VM for this purpose.



NOTE: Before you can update the vGW Security Design VM offline, you must obtain the update ISO from the Juniper Networks Support team and mount it on the vGW Security Design VM.

Before you can update the vGW Security Design VM offline, you must obtain the update ISO from the Juniper Networks Support team and mount it on the vGW Security Design VM.

To perform a manual update offline:

1. In the Settings module Appliance Settings > Updates > section, click **Advanced**.
2. Obtain the update ISO from the Juniper Networks support team, and mount it on the vGW Security Design VM.
3. Select the **Offline Update** check box.
4. Click **Connect Update Media**.
5. To check for updates first, click **Check for Updates**.
6. Click **Update Now**.

Updating the vGW Security Design VM Offline

This procedure explains how to update the vGW Security Design VM offline. Before you can update the vGW Security Design VM offline, you must obtain the update ISO from Juniper Networks Support and mount it on the vGW Security Design VM.

To perform a manual update offline:

1. In the Settings module Appliance Settings > Updates > section, click **Advanced**.
2. Select the **Offline Update** check box.
3. Click **Connect Update Media**.

In the Appliance Settings > Updates > Security Design vGW Update section, click **Check for Updates**.

4. To check for updates first, click **Check for Updates**.
5. Click **Update Now**.

Related Documentation

- [Understanding the vGW Series Update Settings](#)
- [Updating Individual vGW Security VMs on page 321](#)
- [Updating vGW Security VMs in Batch Mode on page 322](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Updating Individual vGW Security VMs

This topic covers how to update an individual vGW Security VM online and offline.

Before you update vGW security VMs, update the vGW Security Design VM. See [“Updating the vGW Security Design VM” on page 319](#)

To update an individual vGW Security VM online, that is, with Internet access:

1. Navigate to the Settings module Security Settings > Security VM Settings page.
The Security VMs pane at the top of the page includes a table with a row for each deployed vGW Security VM. Among other information, the table shows the version of the vGW Security VM.
2. To display a pane that allows you to see detailed configuration information for an individual vGW Security VM, click the row for that vGW Security VM.
3. Select the Updates tab.
4. To check for Updates first to determine if there is a version beyond the one that is currently installed, click **Check for Updates**.
5. If an update exists, click **Update Now**.

This procedure explains how to update the vGW Security VM offline, that is, without using Internet access. Before you can update the vGW Security VM offline, you must obtain the update ISO from the Juniper Networks Support team. To perform a manual update offline:

For SVM migration

1. 1. In VMware attach cd/dvd pointing to the ISO 2. Go to 'Settings->Updates->Advanced->' Select 'Migrate and Offline Update' (In the migrate 'vGW Product Version' box enter "vnf.altornetworks.com@altor:vf-fullers-beta[~vmsafe]")
2. 3. Mount the media by selecting 'Connect Update Media' then select 'Check for Updates' you should see 5.5 listed and then you can proceed with the update



NOTE: If you encounter problems, because, for example, VMs are not moved out correctly though vMotion, on the Security VM Settings page in the Settings module, click **Updates**.

This procedure explains how to update the vGW Security VM offline, that is, without using Internet access. Before you can update the vGW Security VM offline, you must obtain the update ISO from the Juniper Networks Support team.

This procedure explains how to update the vGW Security VM offline, that is, without using Internet access. Before you can update the vGW Security VM offline, you must obtain the update ISO from the Juniper Networks Support team. To perform a manual update offline:

1. Navigate to the Appliance Settings Updates section in the Settings module.
2. Click **Updates**.
3. Click **Advanced**.
4. To check for updates first, click **Check for Updates**.
5. Select the **Offline Update** check box.

To enable offline updates (without Internet access), you must use an ISO image for the vGW Security VM.

6. Click **Update Now**.

**Related
Documentation**

- [Understanding the vGW Series Update Settings](#)
- [Updating the vGW Security Design VM on page 319](#)
- [Updating vGW Security VMs in Batch Mode on page 322](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Updating vGW Security VMs in Batch Mode

This topic explains how to update vGW Security VMs as a group in batch mode.

When you update vGW Security VMs in batch mode, you can run the updates immediately or schedule them to run later.

To set up the system to update vGW Security VMs in batch mode:

1. Navigate to the Settings module Appliance Settings > Updates page.
2. In the Security VM Batch Updates pane, enter the **Custom Product** version.
3. Select the check boxes for the vGW Security VMs that you want to update.
You can select all of them at once.

4. Specify whether you want the updates to run when the ESX/ESXi host is in Maintenance mode. Select:
 - **Always**, in which case logs are not lost.
 - **As needed** for kernel driver updates only.
 - **Never**.
5. Using the **Start Time** option buttons, specify when to run the batch update process.
 - To begin the batch update process immediately, select **Now**. Click **Update**.
 - To schedule the batch update, select **Later**.
 - a. Enter a start date and a start time.
 - b. Optionally, enter an end time.

If you specify an end time, vGW Series completes any update that is in progress when the end time is reached. However, it will not begin any new vGW Security VM updates.
 - c. Optionally, enter an e-mail account to which an update status message is sent when the update either completes or is interrupted.

If you specify an e-mail address for Status email, a message reporting on the vGW Security VMs that were updated and those that are pending is sent to the recipient.

**Related
Documentation**

- [Understanding the vGW Series Update Settings](#)
- [Updating the vGW Security Design VM on page 319](#)
- [Updating Individual vGW Security VMs on page 321](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Configuring the vGW Series Network Settings

This topic covers the Settings module Appliance Settings > Network Settings > Network Configuration page that allows you to change the name of the vGW Security Design VM, the default DNS settings, and the IPv4 or IPv6 default address parameters that are set during installation. It explains how to configure the vGW Security Design VM not to use dual stack.

- [The Network Configuration Page on page 324](#)
- [Changing the Host Name and DNS Settings on page 324](#)
- [Configuring Addresses for the vGW Security Design VM Interface for Communication With vGW Security VMs on page 325](#)

The Network Configuration Page

The vGW Security Design VM uses its Interface 1 virtual NIC (vNIC) for management communication with vGW Security VMs. This interface must be reachable by the management vNICs of all vGW Security VMs. By default, the vGW Security Design VM's Interface 1 is configured for dual stack with DHCP configured to acquire its IPv4 address and DHCPv6 configured to acquire its IPv6 address. [Figure 147 on page 324](#) shows the Network Configuration page that you can use to change these values.

Figure 147: Network Configuration Settings

The screenshot displays the Juniper Networks vGW Security Design VM Network Configuration page. The interface includes a top navigation bar with icons for Main, Network, Firewall, IDS, AntiVirus, Intrusion, Compliance, Reports, Settings, and a status screen. A left sidebar lists 'vGW Application Settings' (Status & License, vCenter Integration, Multi-Center, Installation, Install Settings, Administrators, Active Directory, Machines, High Availability, E-Mail and Reporting, Registry Values) and 'Security Settings' (Global, Security VM Settings, AV Settings, IDS Settings, IDS Signatures, Alerting). The main content area is titled 'Network Configuration' and contains a warning message: 'Notice! Changing Network Settings is not recommended while performing any configuration action that interacts with the vCenter, such as install, uninstall, or update of the Security Design vGW or a firewall.' Below this, the 'Host Name' is set to 'Security_Design_vGW'. The 'DNS Settings' section has 'Use DHCP to Get DNS' checked, with a Primary DNS Server of '172.24.80.10' and a Search Domain. The 'Interface 1' section shows 'IPv4' set to 'DHCP' and 'IPv6' set to 'DHCPv6'. The IP Address is '10.159.24.222', Netmask is '255.255.255.0', Default Gateway is '10.159.24.1', and MAC Address is '00:50:56:90:28:49'. A 'Save' button is at the bottom right.

You can change the default configuration in these ways:

- You can change the vGW Security Design VM and how it gains access to a DNS server.
- You can change the way that the vGW Security Design VM acquires its IPv4 and IPv6 addresses.
- You can configure the vGW Security Design VM not to use dual stack.



WARNING: Do not change the Network Settings during any configuration that involves vGW Security Design VM interaction with VMware vCenter. This includes installing, un-installing, or updating the vGW Security Design VM or a firewall (vGW Security VM).

Changing the Host Name and DNS Settings

You can change the name of the vGW Security Design VM and the default DNS settings using the following sections and their fields on the Network Configuration page.

- **Host Name:** This field allows you to change the name of the vGW Security Design VM management center, `Security_Design_vGW`, that was set by default during installation.

- **DNS Settings:** You can configure the vGW Security Design VM to use either of the following methods to obtain IP address of the Domain Name System (DNS) server to be used:
 - **Use DHCP to Get DNS:** If you want to use Dynamic Host Configuration Protocol (DHCP) to get the IP address of a DNS server dynamically select this option.
By default, the vGW Security Design VM is configured to use this method.
 - **Primary DNS Server:** To use a particular DNS server, de-select **Use DHCP to Get DNS**. Then specify the IP address of the primary DNS server, and optionally, a secondary one.
 - **Search Domain:** You can specify a search domain to use for resolving system names and addresses within vGW Security Design VM reports. To specify more than one search domain, use spaces to separate the domain specifications.

Configuring Addresses for the vGW Security Design VM Interface for Communication With vGW Security VMs

By default, the vGW Security Design VM's Interface 1 is configured for dual stack support with DHCP configured to acquire its IPv4 address and DHCPv6 configured to acquire its IPv6 address.

This section covers how to change the default IP address parameters configured for Interface 1.

- [Changing the Way vGW Security Design VM Acquires Its Interface 1 IP Addresses on page 325](#)
- [Configuring the vGW Security Design VM Not to Use Dual Stack on page 326](#)

Changing the Way vGW Security Design VM Acquires Its Interface 1 IP Addresses

Select how you want vGW Security Design VM to acquire its IPv4 and IPv6 addresses from the lists associated with the following fields:

- **IPv4:**

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

- **DHCP**

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

- **Static IP**

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

- **IPv6:**

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

- **DHCPv6**

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

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

- **Autoconfiguration**

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

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

- **Static IP**

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



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.

Configuring the vGW Security Design VM Not to Use Dual Stack

By default, the vGW Security Design VM is configured for dual stack so that it can communicate with vGW Security VMs that have either IPv4 or IPv6 addresses. You can change the configuration causing it to use either IPv4 addressing or IPv6 addressing alone for communication with vGW Security VMs.

Use the following fields in the Network Configuration Interface 1 pane to cause the vGW Security VM to use a single IP address:

- To use only IPv4 for vGW Security Design VM management communication with its vGW Security VMs, disable IPv6. On the displayed list for the **IPv6:** box, select **Disabled**.

- To use only IPv6 for vGW Security Design VM management communication with vGW Security VMs, disable IPv4. On the displayed list for the **IPv4:** box, select **Disabled**.

In an environment in which the vGW Security Design VM is configured for dual stack communication between the vGW Security Design VM and vGW Security VMs, problems should not exist. Some vGW Security VMs might have IPv4 addresses while others have IPv6 addresses. The environment might also include a standby, or secondary, vGW Security Design VM used for high availability with either type of IP address and that, too, would pose no problems with a dual stack vGW Security Design VM. The vGW Security Design VM can communicate using either protocol.

In environments in which vGW Security VMs and the vGW Security Design VM standby device are configured for dual stack and the primary vGW Security Design VM is not, communication problems should also not exist. Regardless of the type of IP address bound to the vGW Security Design VM's management interface, it would be able to communicate with the management interface of the vGW Security VM or the standby device using their IP address of the same protocol family type.

However, problems will occur if you change the dual stack configuration for the vGW Security Design VM so that it has only one IP address assigned to its Interface 1 vNIC and the management interfaces of the vGW Security VMs and the standby vGW Security VM are configured with only one IP address whose type differs from that of the vGW Security Design VM. For example, if you change the configuration so that the vGW Security Design VM's Interface 1 has only an IPv6 assigned to it, communication problems with any vGW Security VMs with IPv4 addresses will occur. That holds true for the standby vGW Security Design VM also, if one was configured and it had an IPv4 address bound to it. It also holds true for a secondary vGW Security VM, if one was configured with a single IP address that differed in type from the single IP address configured for the management interface of the vGW Security Design VM with which it was intended to communicate.

In circumstances where the IP address types differ, vGW Series presents the following error messages:

- When your environment includes a vGW Security VM—called SVM1 for example—that has only an IPv6 address bound to it, if you attempt to change the vGW Security Design VM from dual stack to single with only an IPv4 address bound to it, vGW Series displays the following message:
 "The interface for management communications must have an IPv6 configuration, because Security VM SVM1 has only IPv6 interface."
- When your environment includes a vGW Security VM—called SVM2 for example—that has only an IPv4 address bound to it, if you attempt to change the vGW Security Design VM from dual stack to single with only an IPv6 address bound to it, vGW Series displays the following message:
 "The interface for management communications must have an IPv4 configuration, because Security VM SVM2 has only IPv4 interface."
- When your environment has a standby vGW Security Design VM that has only an IPv6 address bound to it, if you attempt to change the vGW Security Design VM from dual

stack to single to single with only an IPv4 address bound to it, vGW Series displays the following message:

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

- Related Documentation**
- [Understanding vGW Series IPv4 and IPv6 Dual Stack Support on page 62](#)
 - [Understanding IPv6 Addressing on page 49](#)
 - [Understanding vGW Series IPv6 Support on page 54](#)
 - [Understanding the vGW Security Design VM on page 67](#)

Configuring vGW Series Proxy Settings

You use the Settings module Appliance Settings > Proxy Settings page to enter information about a proxy server, if one is required to make outbound http/https connections. The Security Design vGW connects to the Juniper Networks update server to check for download software updates. If this VM does not have direct access to the Internet, a proxy can be used. All update communications uses HTTPS.

You enter the IP address, port, and user credentials for the proxy server on the Proxy Settings page. vGW Series sends HTTPS (TCP 443) requests to Juniper Networks vGW Internet update servers to pull the latest available software. vGW Series supports both IPv4 and IPv6 addresses for proxy servers.

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

Configuring vGW Series Time Settings

You use the Settings module Appliance Settings > Time Settings page to specify the time zone and current time settings crucial to the proper operation of vGW Series and to specify settings for NTP servers.

It is essential that the vGW Security Design VM have the correct time zone and that it has access to an NTP server. All system logs, security logs, security policy deployment, and other data are time-stamped. If the time setting is not correct, these data will be marked with the wrong time. vGW Security VMs installed on ESX/ESXi hosts synchronize their time settings with that configured on the vGW Security Design vGW.

If you do not have an internal NTP server, you can use the preconfigured NTP servers or another Internet-based NTP server. vGW Series supports both IPv4 and IPv6 addressing for NTP servers.

- Related Documentation**
- [Understanding vGW Series on page 3](#)
 - [Configuring vGW Series Proxy Settings on page 328](#)
 - [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 349](#)

Understanding the vGW Series Backup and Restore Feature

Network and security groups at many companies typically backup and restore configurations for their hardware device systems. In fact for many organizations configuration backup is part of required configuration management practices.

To address this requirement for virtualized devices, vGW Series includes a feature that allows you to back up your vGW Security Design VM configuration to a file store or locally. When necessary, you can easily restore one of your backup versions.

You can run a modified version of the installation wizard that allows you to skip configuration settings that are backed up. After you run the installation wizard and log into the vGW Security Design VM, you can specify the vGW Security Design VM backup configuration to use from your backup location easily.

For details on how to configure settings for the backup and restore feature, see [“Configuring the vGW Series Backup and Restore Feature” on page 331](#).

If the vGW Security Design Center VM configuration that was backed up used a static IP address, the restored version of it has the same IP address. In this case, agents can begin to communicate with the restored vGW Security Design VM immediately after it is started. vGW Series supports both IPv4 and IPv6 addresses.

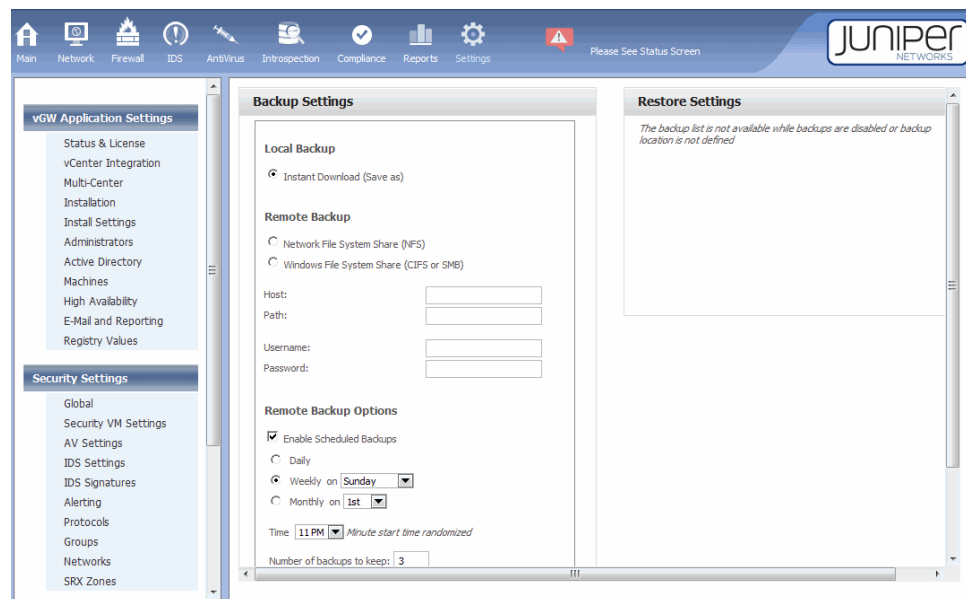
The issue is you can't use the 'Windows File System Share (CIFS or SMB)' option if you are using only IPv6. This option requires IPv4



NOTE: You cannot use the Windows File System Share (CIFS or SMB) option if you are using only IPv6. This option requires IPv4.

See [Figure 148 on page 330](#).

Figure 148: Settings Module Backup and Restore Settings



vGW Series backs up and restores the following content:

- All configuration information that you configured using the vGW Security Design VM, including:
 - machines
 - networks
 - groups
 - protocols
 - security policies objects
 - policies associated with groups
 - Smart Groups group membership and logic
 - Static Groups VM membership. Note that the VM-ID/UUID will change during the backup and restore process.
 - Administrator accounts. Administrator passwords are exported in a safe and secure manner.



NOTE: All source tables are backed up except those for connections, alerts, and idp_alerts. You must use another tool or process that copies the entire VM (vmdk) to back up that information.

The vGW Security Design VM backup-and-restore feature allows you to:

- Specify where to back up the files, locally or to a remote store.
- Specify the number of backup files to retain. You can remove all backup copies whenever you choose to.
- From among the backed-up versions, select the configuration to restore.
- Schedule when to back up the configuration.

Related Documentation

- [Configuring the vGW Series Backup and Restore Feature on page 331](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Configuring the vGW Series Backup and Restore Feature

vGW Series provides a backup and restore feature that you can use to create and store multiple backups of your vGW Security Design VM configuration and easily restore a backed up version. This topic explains how to configure settings for it. It also explains what you must do after you restore the vGW Security Design VM from a backup version. [Figure 149 on page 331](#) shows the Settings module Appliance Settings > System Settings Backup page that you use for this purpose.

Figure 149: Settings Module Backup and Restore Settings

The screenshot shows the Juniper Networks vGW Series Settings Module. The left sidebar contains a tree view with categories like vGW Application Settings and Security Settings. The main content area is titled 'Backup Settings' and is divided into two panes. The left pane, 'Local Backup', has a radio button for 'Instant Download (Save as)'. The right pane, 'Remote Backup', has radio buttons for 'Network File System Share (NFS)' and 'Windows File System Share (CIFS or SMB)', with input fields for Host, Path, Username, and Password. Below these are 'Remote Backup Options' including 'Enable Scheduled Backups' (checked), frequency (Weekly on Sunday), time (11 PM), and 'Number of backups to keep' (3). The rightmost pane, 'Restore Settings', is empty and shows a message: 'The backup list is not available while backups are disabled or backup location is not defined'.

You can configure your system to back up the vGW Security Design VM in either of the following ways:

- **Local Backup:** This option backs up the vGW Security configuration immediately as an instant download, similar to a Save As function.

- **Remote Backup:** This option allows you to back up the vGW Security Design VM remotely at a scheduled time. You can also use it to back up the vGW Security Design VM now.



NOTE: You cannot back up the vGW Security Design VM over IPv6 networks.

To back up the vGW Security Design VM remotely:

1. In the **Backup Settings** pane under **Remote Backup**, specify where you want the backup to be stored. Select:
 - Network File System Share (NFS)
 - Windows File System Share (CIFS or SMB)

If you want to use the Windows Files System (CIFS (Common Internet File System) or SMB (Server Message Block) as the file store, specify the following information:

 - The name of the host and its path.
 - The username and password used to access the file share.
2. Configure the **Remote Backup Options**.
 - a. To schedule the backup, select **Enable Scheduled Backups**.
 - Specify the date and time:
 - For **Daily**, select it.
 - For **Weekly**, select the day.
 - For **Monthly**, select the date of the month.
 - In any of these cases, in the **Time** box, select the time when the backup should begin.
 - b. In the **Number of backups to keep**:, specify the number of backup versions to create and write to the file share.
3. Select **Backup before software update** to direct the vGW Series to back up the vGW Security Design VM before it is updated, whether the update is automatic or manual.
4. Click **Save** to save your backup configuration definition.
5. To back up your configuration immediately, click **Backup Now**.

The Restore Settings pane allows you to restore a backup file from the location where you stored the files.



NOTE: Clicking Restore causes the system to be restarted.

- To upload a local restore file, in the Local Restore section of the Restore Settings pane:
 1. Click **Browse...** to locate the backup file.

To clear the file selection, press **Clear**. Pressing Clear does not delete the backup file. It simply clears the browse box to allow you to select a different file.

2. Click **Restore**.

- To restore a backup file from a remote location, in the Remote Restore section of the Restore Settings pane:

1. Select a backup file from the list. The list identifies the name of the file, the date it was created, and the file size.

To refresh the Remote Restore configuration backup list, click **Retrieve File List**.

2. Click **Restore**.



WARNING: After you restore the vGW Security Design VM, you must reconfigure high availability (HA) for it. See the following procedure.

After you restore the vGW Security Design VM:

1. Verify that it is working properly.
2. Cancel the standby (secondary HA) vGW Security Design VM.
3. Reconfigure HA for the restored vGW Security Design VM.

For details on how to reconfigure HA, see “[Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability](#)” on page 349.

Related Documentation

- [Understanding the vGW Series Backup and Restore Feature on page 329](#)
- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Series Settings Module on page 203](#)

Understanding vGW Series Log Collection

vGW Series Collection Tool allows you to generate log collections for the vGW Security Design VM and vGW Security VMs. You can generate log collections for:

- Only the vGW Security Design VM.
- One vGW Security VM.
- The vGW Security Design VM and one or more vGW Security VMs collectively from the same point of access.

In this case, you select the vGW Security VMs along with the vGW Security Design VM, and log collections will be generated for them also.

You can also generate log collections for a secondary vGW Security VM if you configured one for a vGW Security VM for high availability.

You can provide the log collection files to the Juniper Networks Support team to be used for diagnostic troubleshooting purposes. The Collection Tool is available on the Settings module Appliance Settings > Log Collection page.

- [Log Collection on page 334](#)
- [Generating the Log Collections on page 334](#)
- [Uploading the File on page 335](#)
- [Downloading the File on page 336](#)
- [Using a Method Other Than the vGW Security Design VM to Generate Log Collections for It on page 336](#)

Log Collection

For some reason you might encounter problems that necessitate troubleshooting the vGW Security Design VM itself and maybe one or more vGW Security VMs that exhibit problems. The Collection Tool allows you to generate information useful in solving the problem.

When you generate log collections for the vGW Security Design VM and vGW Security VMs together, they are included in a single compressed archive file. The (TGZ) zip file contains a separate zip for each component—a zip file for the vGW Security Design VM log collection and separate zip files for the log collections of each vGW Security VM that you selected. If no vGW Security VMs were selected, the zip file contains only the logs collected for the vGW Security Design VM.



NOTE: Because vGW Security VM log collections are copied to the vGW Security Design VM, it is recommended that you not initiate log collections for more than three vGW Security VMs concurrently. Although the recommendation is not a restriction, typically problems encountered are constrained to a small number of vGW Security VMs, so it should not be necessary to include a large number vGW Security VMs.

To remind you of the recommendation, vGW Series displays a confirmation alert after you have selected more than three vGW Security Design VMs.

After the log collections have completed, you can directly upload the file to a support server or you can download a copy of the file to submit to the Juniper support team manually. Juniper Networks recommends that you upload the file to the support server.

Generating the Log Collections

If you do not select the check box for any of the vGW Security VMs, vGW Series generates log collections only for the vGW Security Design VM. To take this action:

1. Click **Start New Collection**.

The resulting zip file contains only the vGW Security Design VM logs.

2. Follow the instructions in [Uploading the File](#) or [Downloading the File](#) to send the log collections to the Juniper Support team.

To use the Collection Tool to generate log collections for the vGW Security Design and one or more vGW Security VMs:

1. Select the **Get Logs** check box to the right of the name of each vGW Security VM for which you want to generate a log collection. The Collection Tool generates relevant log and system files and it compresses them in a TGZ zip file. The log collections are packaged together in a single zip file with the logs for the vGW Security Design VM.
2. Click **Start New Collection**.

The log collection process begins. The Collection Tool generates the log collection for all entities in parallel. When all the log collections are copied to the vGW Security Design VM, it shows a message of **Done! The log is now ready**.

You can either upload the file to the support center or download a copy of the zip file containing all the log collections to your local system to submit to the Juniper Support team manually.



NOTE: Juniper Networks recommends that you upload the file to the support server.

Uploading the File

When you click Upload in the Upload Log Collection pane, the newest zip file is uploaded, and an ID is returned to you. The upload process encrypts the file (through AES-256), and it transfers it to a protected server. Before you upload the file, briefly describe the problem in the comment field.

To upload the file:

1. Provide a brief description of the problem in the scroll box in the **Upload Log Collection** pane, and any other comments that you would like to submit.
2. Click **Upload**.

A **Submitting collection to support server....** progress message is reported. When the process completes, the message **Submission successful. Upload ID *id-number*** is displayed.

Make note of the ID, which you use to track your submission. You should refer to this ID in trouble tickets or other communication with the Juniper Networks support team about the problem.

Downloading the File

If you click **Download** in the **Download Log Collection** pane to download the file, you can send the file to Juniper Networks Support through e-mail at any time, or you can post the log collection to a server.

The downloaded file is called `datacollection-date.tar.gz`.

Using a Method Other Than the vGW Security Design VM to Generate Log Collections for It

When it is not possible to generate a log collection from the vGW Security Design VM, you can run the Collection Tool using another method. You can use this method to generate a log collection only for the vGW Security Design VM.

From the vGW command line interface (vGW CLI), you can use the `logs collect` command to generate log collections for the vGW Security Design. The command prints out the location of the file. You can then copy the file to another location using `scp`.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Viewing the vGW Series Logs on page 336](#)
- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)

Viewing the vGW Series Logs

You can use the Settings module Appliance Settings > Log Viewer to view system and application logs for basic system activity monitoring and troubleshooting. You can also select the number of lines that are displayed in the viewer.

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series Log Collection on page 333](#)
- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)

Understanding vGW Series Support Settings

You use the support section of the Settings module in the vGW Security Design VM to:

- reboot the vGW Security Design VM.
- restart vGW Series services.
- enable or disable debugging flags used for troubleshooting.

If you enable debug flags, return to this page after the log files are collected, and click **Debugging OFF**. When the debug setting is enabled, many log files are generated which could cause disk space usage problems.

Click Advanced in the Debug Flags pane to display a list of debug flags that you can set. For example, you can enable the active.directory flag to generate additional troubleshooting information in /usr/lib/tomcat/webapps/ROOT/log/debug.log.0.

- enable or disable SSH remote access to the vGW Security Design VM.

When you enable SSH, you can administer the vGW Series through an SSH client, such as PUTTY. This allows security teams to access the vGW command line of the vGW Security Design VM and the vGW Security VM components without having to use the vSphere Client.

When you access the vGW Security Design VM or vGW Security VM(s) through SSH, you are presented with a command-line interface. The command-line interface supports a variety of system options. Enter ? or enter **help** at the command-line prompt to view a list of supported vGW Series commands.

**Related
Documentation**

- [Understanding vGW Series on page 3](#)
- [Understanding vGW Series Log Collection on page 333](#)

vGW Series Status Alerts

- [Understanding vGW Series Status and Alerts on page 339](#)

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 339](#)
- [Alerts on page 339](#)
- [E-Mail Alert Settings on page 340](#)
- [SNMP Trap Settings on page 340](#)
- [AutoConfig and Multicast Alerts on page 340](#)

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](#)

High Availability and Fault Tolerance

- [Understanding the vGW Series High Availability Solution on page 343](#)
- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 349](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)
- [Understanding vGW Series Fault Tolerance Support on page 354](#)

Understanding the vGW Series High Availability Solution

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

- [vGW Series HA on page 343](#)
- [vGW Series HA and VMware HA on page 344](#)
- [vGW Series HA for the vGW Security Design VM on page 344](#)
- [vGW Security Design VM HA Behavior on page 346](#)
- [vGW Series HA for the vGW Security VM on page 348](#)

vGW Series HA

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



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

vGW Series HA and VMware HA

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



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

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

vGW Series HA for the vGW Security Design VM

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

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



NOTE: vGW Series high availability is meant to be used as an emergency solution, not as a replacement system. If the primary vGW Security Design VM fails, it can be recovered from a backup or snapshot copy. For details, see [“Configuring the vGW Series Backup and Restore Feature” on page 331](#).

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

[“Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability” on page 349](#) explains the process for creating a secondary vGW Security Design VM.

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

- **IPv4:**

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

- **DHCP**

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

- **Static IP**

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

- **IPv6:**

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

- **DHCPv6**

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

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

- **Autoconfiguration**

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

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

- **Static IP**

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

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

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

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

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

See ["Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability"](#) on page 349.



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

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

By default, this parameter is set to true.

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

vGW Security Design VM HA Behavior

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

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

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

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



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

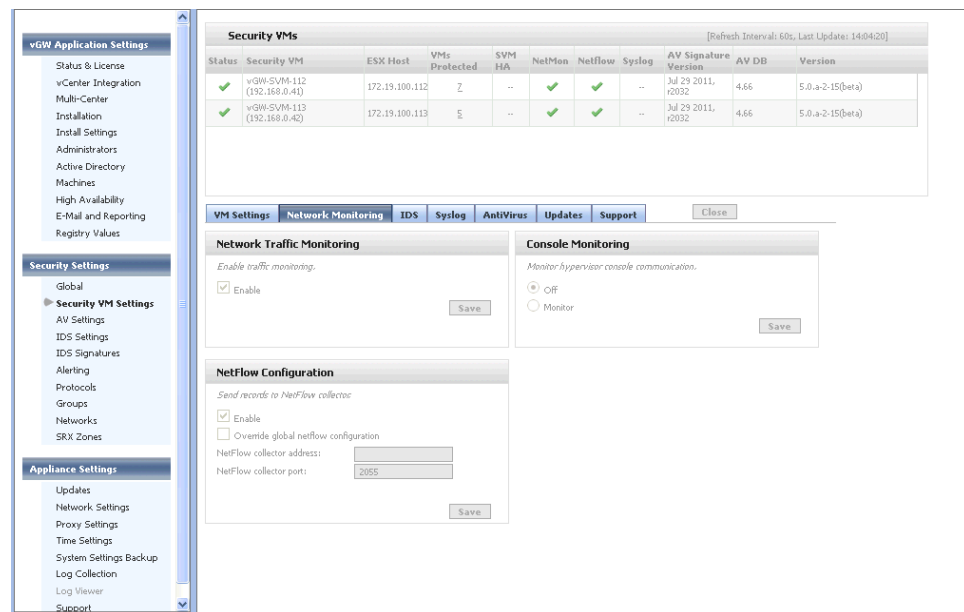
- vGW Series does not synchronize events back from the secondary vGW Security Design VM to the primary one.
 - You cannot create compliance rules.
 - Changes to Network Monitoring and NetFlow, IDS, and AntiVirus events and changes to statistics are not viewable unless you configure NetFlow and Syslog from the vGW Security VM for individual vGW Security VMs.



NOTE: Although you cannot view their activity from the secondary vGW Security Design VM, these features continue to work when the primary vGW Security Design VM is unavailable.

Figure 150 on page 348 shows the Settings module page that you use to configure Network Monitoring and NetFlow for individual vGW Security VMs.

Figure 150: Configuring Network Monitoring and NetFlow Settings



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

vGW Series HA for the vGW Security VM

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

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

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

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

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

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

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

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

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

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

For details on how to install a vGW Security VM, see [“Installing a Secondary vGW Security VM for High Availability” on page 353](#).

It is important to consider that the IP protocol address type of the IP address bound to the management interface of the secondary vGW Security VM must correspond to that of the vGW Security Design VM management interface with which it communicates. However, if both or either one is configured for dual stack, communication problems should not occur. If both are not configured for dual stack and the types of the IP addresses bound to their management interfaces differs, communication problems will ensue. For further information, see [“Installing vGW Security VMs on ESX/ESXi Hosts” on page 218](#).

This pane allows you to change the IP protocol family that is used for the vGW Security VM management interface when that protocol does not match that of the vGW Security Design VM with which it must communicate. For information on conditions that would cause an IP address type mismatch between the management interfaces of the vGW Security VM and the vGW Security Design VM, see [“Setting Up vGW Series” on page 36](#) and [“Installing vGW Security VMs on ESX/ESXi Hosts” on page 218](#).

Related Documentation

- [Understanding vGW Series on page 3](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding the vGW Security VM on page 68](#)
- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 349](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353](#)
- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)

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

This topic explains how to install an additional vGW Security Design VM to be used when the primary one is unavailable. You can install more than one additional vGW Security

Design VM. It also explains how to configure the primary vGW Security Design VM for HA and how to determine the secondary one to use for it. The process entails:

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



CAUTION: Be sure to back up your primary vGW Security Design VM. vGW Series does not rebuild a primary vGW Security Design VM from a secondary one created for HA. For details on backing up the primary vGW Security Design VM, see [“Configuring the vGW Series Backup and Restore Feature” on page 331](#).

To create a secondary vGW Security Design VM:

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



NOTE: If you need further information about installing the secondary vGW Security Design VM, you can read about how it is done for the primary vGW Security Design VM. See [“Understanding the Open Virtualization Format OVA Template Method” on page 23](#) and [“Using the OVA Single File Method to Integrate the vGW Security Design VM with VMware” on page 34](#), and related topics that they refer to.

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

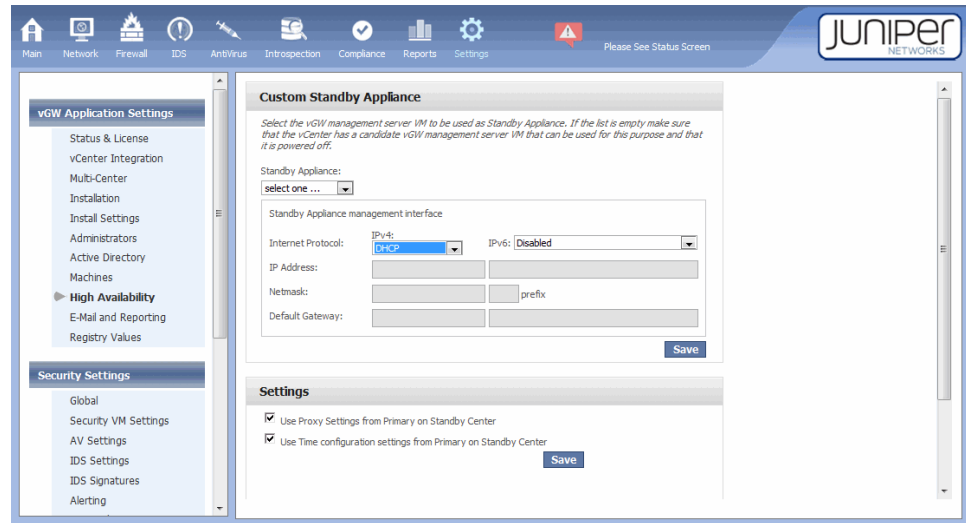


CAUTION: After the import completes, do not power on the newly created secondary vGW Security Design VM.

To configure the primary vGW Security Design VM for HA:

1. Configure the vGW Security Design VM for HA in the Settings module:
 - a. To configure the secondary vGW Security Design VM, select **vGW Application Settings > High Availability**. See [Figure 151 on page 351](#).

Figure 151: Configuring the Secondary vGW Security Design VM



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



NOTE: IPv4 DHCP is enabled by default.

From the Internet Protocol list select:

- For IPv4
 - **Disabled**
Disable IPv4 and use an IPv6 address for the secondary vGW Security Design VM.
 - **DHCP**
Use DHCP to assign an IPv4 address dynamically to the secondary vGW Security Design VM.
 - **Static IP**
Specify a static IPv4 address, its network mask routing prefix, and the default gateway to use for the secondary vGW Security Design VM.
- For IPv6:

- **Disabled**

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

- **DHCPv6**

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

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

- **Autoconfiguration**

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

- **Static IP**

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

- d. Click **Save**.

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

- a. Specify whether to use the proxy settings configured for the primary vGW Security Design VM for the standby (secondary) one. See ["Configuring vGW Series Proxy Settings" on page 328](#).

The Security Design vGW connects to the Juniper Networks update server to check for available downloads of software updates. If the server does not have direct access to the Internet, a proxy can be used. For the primary vGW Security Design VM, the Settings module Appliance Settings > Proxy Settings page specifies configuration information about a proxy server, if one is required to make outbound http/https connections.

- b. Specify whether to use the time configuration settings configured for the primary vGW Security Design VM on the standby (secondary) one. See ["Configuring vGW Series Time Settings" on page 328](#).

- c. Click **Save**.

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

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

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

Related Documentation

- [Understanding the vGW Series High Availability Solution on page 343](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Installing a Secondary vGW Security VM for High Availability on page 353.](#)
- [Preparing to Integrate the vGW Series with the VMware Environment on page 21](#)
- [Understanding vGW Series Fault Tolerance Support on page 354](#)

Installing a Secondary vGW Security VM for High Availability

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

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

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

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

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

To clone the existing, primary vGW Security VM:

1. Select Settings > Security Settings > Security VM Settings.
2. Click the row for the vGW Security VM that you want to duplicate.
3. In the High Availability pane, click **Configure**.

4. Enter information for the secondary vGW Security VM. Specify the appropriate IP address information, management network, and data store location.
5. Click **Configure**.



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

Related Documentation

- [Installing an Additional vGW Security Design VM and Configuring the Primary vGW Security Design VM to Use It for High Availability on page 349](#)
- [Understanding the vGW Series High Availability Solution on page 343](#)

Understanding vGW Series Fault Tolerance Support

This topic contains the following sections:

- [About vGW Series Fault-Tolerance on page 354](#)
- [vGW Series Fault Tolerance in the vGW Series on page 355](#)
- [Enabling Fault Tolerance for a Virtual Machine on page 355](#)

About vGW Series Fault-Tolerance

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

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

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

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

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

vGW Series Fault Tolerance in the vGW Series

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

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

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

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

- You cannot define a policy for the secondary VM.

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

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

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

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

Enabling Fault Tolerance for a Virtual Machine

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

To enable FT for a VM in the vCenter:

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

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

You c

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

**Related
Documentation**

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

PART 3

Juniper Networks Products Interoperability

This part includes the following chapters:

- [vGW Series Interoperability with Juniper Networks Products on page 359](#)

CHAPTER 20

vGW Series Interoperability with Juniper Networks Products

This chapter covers the following topics:

For additional information on vGW Series integration with other Juniper Networks products, including in-depth coverage of STRM Series, SRX Series for zone synchronization, and SRX Series and IDP Series, see the Security Virtualization Application note at <http://www.juniper.net/us/en/solutions/enterprise/data-center/secure>.

- [vGW Series and SRX Series Security Zones on page 359](#)
- [Enabling the Junoscript Interface for vGW Series on page 360](#)
- [Configuring Zone Objects for vGW Series Interoperability with SRX Series Devices on page 362](#)
- [About Populating vGW Series Records to SRX Series Zone Address Books on page 363](#)
- [Validating vGW Series Interoperability with SRX Series Zones on page 363](#)
- [Configuring vGW Series to Send Syslog and Netflow Data to Juniper Networks STRM Series Devices on page 364](#)
- [Configuring the vGW Series and IDP Series Inter-Operation on page 367](#)

vGW Series and SRX Series Security Zones

This topic includes the following sections:

- [About SRX Series Services Gateway Security Zones on page 359](#)
- [SRX Series Services Gateway Zones and the vGW Series on page 360](#)

About SRX Series Services Gateway Security Zones

A security zone is a collection of one or more network segments on SRX Series devices requiring the regulation of inbound and outbound traffic through policies.

Security zones are logical entities to which one or more interfaces on the SRX Series device are bound.

On a single SRX Series device, you can configure multiple security zones, dividing the network into segments to which you can apply various security options to satisfy the

needs of each segment. You can define many security zones, bringing finer granularity to your physical network security design—and without deploying multiple security appliances to do so.

From the perspective of security policies, traffic enters into one security zone and goes out on another security zone. This combination of a from-zone and a to-zone is defined as a context. Each context contains an ordered list of policies.

SRX Series devices support many types of security zones.

SRX Series Services Gateway Zones and the vGW Series

vGW Series zone synchronization feature provides an automated way to link the vGW Series virtualized security layer with the SRX Series Services Gateway physical device and network security.

vGW Series zone feature simplifies VM-to-zone mapping by importing into the virtualized environment zones configured on SRX Series devices.

You can use these zone assignments to:

- Apply zone policies to use between VMs.
- Integrate zones with compliance checking to ensure that VMs are attached only to authorized zones.

The process that the vGW Series undertakes to synchronize SRX Series zones with VMs consists of a number of steps, including defining:

- An SRX object. This process entails obtaining zone configuration information from the SRX Series device, mapping zones to the vGW Series interface, and associating VLANs or network ranges with each zone.
- Zones as Smart Groups within the vGW Series based on the VLANs and the networks associated with each zone.

vGW Series also validates that Smart Groups dynamically associated with each VM are associated with the appropriate zone. This process allows for policy enforcement between vGW Series VMs and SRX Series zone compliance validations.

Related Documentation

- [Understanding vGW Series on page 3](#)
- For additional information on vGW Series integration with other Juniper Networks products, including in-depth coverage of STRM, SRX for zone synchronization, and SRX-IDP, see the Security Virtualization Application note at <http://www.juniper.net/us/en/solutions/enterprise/data-center/secure>.

Enabling the Junoscript Interface for vGW Series

To allow the vGW Series to gain access to the SRX Series device for zone synchronization, you must enable the secure Junoscript XML scripting API. To do so:

1. Generate a digital Secure Sockets Layer (SSL) certificate, and install it on the SRX Series device.
 - a. Enter the following openssl command in your SSH command-line interface on a BSD or Linux system on which openssl is installed. The openssl command generates a self-signed SSL certificate in privacy-enhanced mail (PEM) format. It writes the certificate and an unencrypted 1024-bit RSA private key to the specified file.

```
% openssl req -x509 -nodes -newkey rsa:1024 -keyout mycert.pem -out mycert.pem
```

- b. Type the appropriate information in the identification form, when prompted. For example, type US for the country name.
 - c. Copy the certificate that you generated from the operating system to the SRX Series device. In this example, the certificate is copied to the /var/tmp/ directory.

```
scp mycert.pem user@host:/var/tmp/
```

- d. Install the mycert.pem SSL certificate on the SRX Series device. Using the CLI, enter the following statement in configuration mode:

```
[edit]
user@host# set security certificates local mycert load-key-file
/var/tmp/mycert.pem
```

2. Enable HTTPS for Web management access at the system level. Specify the SSL certificate and the web management port.

You can enable HTTPS access on specified interfaces. If you do not specify an interface, HTTPS is enabled on all interfaces. In this example, ge-0/0/0.0 is used.

```
[edit]
user@host# set system services web-management https local-certificate mycert
user@host# set system services web-management https interface ge-0/0/0.0
user@host# set system services web-management https port 443
```

3. Configure the *zone* to allow HTTPS as the protocol for host inbound traffic for Web management on all of its interfaces.

```
[edit]
user@host# set security zones security-zone trust host-inbound-traffic system services
https
```

4. Configure the IP address for the interface, if it is not already configured.
5. Enable Junoscript communications using the newly created certificate:

```
[edit] user@srx# set system services xnm-ssl local-certificate mycert
```

Related Documentation

- [vGW Series and SRX Series Security Zones on page 359](#)
- [Understanding the vGW Series SRX Zones Settings on page 317](#)
- [About Populating vGW Series Records to SRX Series Zone Address Books on page 363](#)
- [Understanding vGW Series on page 3](#)

Configuring Zone Objects for vGW Series Interoperability with SRX Series Devices

To create a new SRX Series zone object, using the vGW Security Design VM interface:

1. Select the Settings module.
 - a. In the Security Settings box on the left pane, select **SRX Zones**.
 - b. Click the **Add** button on the lower right side of the page.

The Add SRX Zone dialog box appears.

- c. Specify the following information for the SRX Series zone in the Add SRX Zone pane:
 - **Name**—A short descriptive name for the SRX Series zone object. This name is used in VM zone labels.
 - **Host**—Device management IP address on the SRX Series device used to connect to the vGW Security Design VM.
 - **Port**—TCP port used to connect to the SRX Series device through the Junoscript interface.
 - **Login and Password**—Credentials used to authenticate to the SRX Series device.

The account for the SRX Series zone object requires read access to the SRX Series device's zones, interface, network, and routing configuration. Optionally, it requires write access to the Address Book for each zone to populate it with VM entries.

If you do not want the system to enter VM objects in the SRX Series device's address book, you do not need to provide write access.

- **VMs associated with this SRX**—This optional parameter specifies the VMs scope. It can be a smart group that defines VMs that are relevant to the SRX Series device.
2. Define synchronization intervals and relevant interfaces, by clicking **Load Zones** after you save the SRX Series zone object definition.

After the zone synchronization process has completed, a list of zones that the vGW Series retrieved appears. You can select the zones to import into the vGW Series as VM zone groupings.

You can configure zone synchronization to automatically poll the SRX Series device for zone updates.

To configure the vGW Series automatic zone synchronization process to control synchronization update, specify the following information:

- **Update Frequency**—How often to query the SRX Series device for updates (interval).

- Relevant Interfaces—Select the SRX Series device interfaces to be monitored by vGW Series. vGW Series discovers any new zones assigned to the relevant interfaces and adds them for monitoring.

SRX Series zones that participate in the synchronization process are automatically created in the vGW Series as VM Smart Groups. A Smart Group is created based on the following parameters:

- VLANs associated with the SRX Series device interface.
- The subnet defined on the SRX Series device interface and routes defined within a zone.

If the zone synchronization configuration includes a VM associated selection, the group you select is included in the Smart Group Definition.

**Related
Documentation**

- [vGW Series and SRX Series Security Zones on page 359](#)
- [About Populating vGW Series Records to SRX Series Zone Address Books on page 363](#)
- [Understanding vGW Series on page 3](#)

About Populating vGW Series Records to SRX Series Zone Address Books

vGW Series to SRX Series zones synchronization feature allows VM records to be populated in the SRX Series address book for the zone that the VM belongs to. This allows the VM-to-zone mapping validation to occur within the context of the SRX Series device management.

When a VM record is added to an SRX Series device's zone address book, it is created with the name of the VM as defined in vCenter. A string is prepended to the name of the VM in its address book entry to indicate that it is an auto-generated VM record. By default, the string "VM-" is used, but you can change the name in the synchronization dialog box. If you change this string, vGW will attempt to update all of your existing entries to use the new string. Your existing entries are not lost or removed.

**Related
Documentation**

- [vGW Series and SRX Series Security Zones on page 359](#)
- [Configuring Zone Objects for vGW Series Interoperability with SRX Series Devices on page 362](#)
- [Validating vGW Series Interoperability with SRX Series Zones on page 363](#)
- [Understanding vGW Series on page 3](#)

Validating vGW Series Interoperability with SRX Series Zones

When the VM zone attachment information is accessible within the vGW Security Design VM, you can incorporate it into the policy automation and compliance checking procedures.

For VMs that do not meet compliance requirements, immediate action can be taken. You can create a noncompliant group and group policy to lock out noncompliant VMs from the network. Any noncompliant VMs are added to this group.

**Related
Documentation**

- [Configuring Zone Objects for vGW Series Interoperability with SRX Series Devices on page 362](#)
- [About Populating vGW Series Records to SRX Series Zone Address Books on page 363](#)
- [vGW Series and SRX Series Security Zones on page 359](#)
- [Understanding vGW Series on page 3](#)

Configuring vGW Series to Send Syslog and Netflow Data to Juniper Networks STRM Series Devices

Integration of vGW Series with Security Threat Response Manager (STRM) Series devices provides for defense-in-depth control in the virtualized server environment. This topic covers vGW Series Syslog and Netflow integration configuration with the Juniper Networks STRM Series device.

vGW Series and STRM Series integration brings STRM Series benefits such as centralized log and event management, network-wide threat detection, and compliance reporting to the virtualized data center. This integration gives you a single-pane, comprehensive, and consistent view of your physical and virtual infrastructure.

vGW Series and STRM Series have two points of integration. vGW Series exports the following information to the STRM Series device:

- Syslog firewall logs and event messages.
- NetFlow statistics on traffic between virtual machines (VMs).

You use the Settings > Global page to configure the vGW Security Design VM to send Syslog logs and events and NetFlow VM traffic information to the STRM Series device. See [Figure 152 on page 365](#).

Figure 152: vGW Series Configuration for Syslog and NetFlow to a STRM Series Device

The figure displays four screenshots of the vGW Series configuration interface:

- External Inspection Devices:** A table with columns for #, Name, and IP Address. Row 1 is filled with '1', 'STRM', and '10.10.10.8'. Rows 2, 3, and 4 are empty. A 'Save' button is at the bottom right.
- Global Settings Rules:** A table with columns for #, Rule, and Allow. Row 1 is '1', 'IPv6 traffic', and an unchecked checkbox. Row 2 is '2', 'Non-IP and non-ARP traffic', and an unchecked checkbox. A 'Save' button is at the bottom right.
- External Logging:** Radio buttons for 'No Syslog', 'Send Syslog from vGW management server', and 'Send Syslog from Firewalls'. The 'Send Syslog from Firewalls' option is selected, with a checked checkbox for 'Send firewall logs to vGW management server'. Below are fields for 'Syslog Server:' (10.10.10.8), 'Syslog Server Port:' (514), and 'Transport Protocol:' (tcp and udp, with udp selected). A 'Save' button is at the bottom right.
- NetFlow Configuration:** A checked checkbox for 'Enable'. Below are fields for 'NetFlow collector address:' (10.10.10.8) and 'NetFlow collector port:' (2055). A 'Save' button is at the bottom right.

Syslog. For Syslog, you configure information on both vGW Series and the STRM Series device:

- You configure information about the STRM Series device on the vGW Series Settings > Global pane.
- The Syslog format is particular to a specific device. Therefore, for STRM to be able to recognize and parse vGW Series incoming messages, you must specify the vGW Series log source on STRM.

Syslog Configuration on vGW Series.

Configure vGW Series for Syslog external logging to the STRM Series device.

1. In the External Inspection Devices pane, enter STRM for the name of the external device and specify the STRM Series device's IP address.
2. In the External Logging pane, select **Send Syslog from Firewalls**. If you want to send the firewall logs to the vGW Security Design VM also, select the check box.
3. To identify the STRM Series device as the Syslog server, specify its IP address in the External Logging pane.
4. Select UDP as the transport protocol.

vGW Series Configuration on STRM.

1. Define vGW Series as the log source in the STRM Series device to identify the Syslogs that you are sending. See [Figure 153 on page 366](#).

Figure 153: STRM Source Log Definition for vGW Series

Add a log source	
Log Source Name	vGW
Log Source Description	My Virtual Gatekeeper
Log Source Type	Juniper vGW
Protocol Configuration	Syslog
Log Source Identifier	10.10.10.10
Enabled	<input checked="" type="checkbox"/>
Credibility	5
Target Event Collector	eventcollector0 :: strm-console
Coalescing Events	<input checked="" type="checkbox"/>
Store Event Payload	<input checked="" type="checkbox"/>
Please select any groups you would like this log source to be a member of:	

NetFlow.

The STRM Series device can listen for NetFlow messages on port 2055 from any device because NetFlow has a standard format. If you specify 2055 for the port on the vGW Series configuration, you do not need to configure NetFlow on the STRM Series device.

In the Settings > Global > NetFlow Configuration pane, configure vGW Series NetFlow to send VM traffic statistics to the STRM Series device. See [Figure 152 on page 365](#).

1. Select the **Enable** check box.
2. Specify the IP address of the NetFlow collector and the destination port to use.



NOTE: The standard specification is UDP port 2055, but other values like 9555 or 9995 are sometimes used. If you use another value, you must configure vGW Series NetFlow information on the STRM Series device.

**Related
Documentation**

- [Understanding the vGW Series Main Module on page 77](#)
- [Understanding the vGW Series Security Alert Settings on page 297](#)
- [Understanding the vGW Security Design VM on page 67](#)
- [Understanding vGW Series on page 3](#)
- [Adding and Editing vGW Series Machines Definitions \(VMware\) on page 277](#)

Configuring the vGW Series and IDP Series Inter-Operation

Juniper Networks IDP Series Intrusion Detection and Prevention Appliances provides features that protect the network from a wide range of attacks. Using stateful intrusion detection and prevention techniques, the IDP Series provides protection against worms, trojans, spyware, keyloggers, and other malware. Its feature set includes stateful signature detection, protocol and anomaly detection, QoS/DiffServ marking, VLAN-aware rules, role-based administration, separation of domains and management activities, IDP Reporter, and traffic pattern profiling.

Before you configure interoperability between the vGW Series and the IDP Series, you must configure the Intrusion Detection System as an external inspection device and configure an appropriate redirection rule for it using the Global section of the Security Settings module.

The External Inspection Devices page allows you to enter the name and IP address of the device to which traffic is sent for further analysis.

To configure the vGW Series and IDP Series inter-operation:

1. Log into the NSM for your environment.
2. Create a Security Policy for the Inter-VM communication:
 - a. In the notification section of the policy, select **Logging**.
 - b. Enable the policy for traffic between any source and destination.
 - c. Set the action to **None**.

You can inspect traffic anomalies between VMs using this security policy.

3. Enable GRE decapsulation support on the IDS device for which you created the security policy.
4. Select **Device Manager** > Security Devices.
5. Select **Sensor Settings** > **Run-Time Parameters**.
6. Select **Enable GRE decapsulation** support.

To verify that you set the parameter correctly, enter the following command on the command line of the IDP Series device:

```
ser@host# scio const -s s0 get sc_gre_decapsulation
```

After you have completed these steps, you can test the configuration. Once the above steps are complete (including the creation of the External Inspection Device and relevant security policy in vGW Security Design VM) you can test the configuration by triggering any attack in the Juniper Networks database.

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

PART 4

Index

- [Index on page 371](#)

Index

A

address books.....	363
administrators	
changing passwords.....	269
defining and adding.....	269
AntiVirus.....	131
overview.....	131
appliance status.....	207
Application Settings.....	206
authentication	
changing passwords for administrators.....	269
Auto Deploy.....	225
VMware.....	227

B

backup and restore.....	329, 331
-------------------------	----------

C

cloud computing.....	11
compliance check procedures	
VM-to-zone mapping attachment.....	363
Compliance module.....	183
configuring basic system parameters.....	36
customer support.....	xxiv
contacting JTAC.....	xxiv

D

database status.....	207
delegate centers.....	244
DNS services.....	15
documentation	
comments on.....	xxiv
dual stack.....	62

E

Enforcer Profiles tab.....	172
ESX/ESXi hosts.....	13
events and alerts.....	77

F

Firewall module.....	93
used with Network module.....	90
firewall policy rules for vGW components.....	120

G

Global Settings	282
Gold Image.....	171, 172
group policy	
for checking VM-to-zone mapping	363
Groups.....	300

H

high availability.....	343
distributed resource schedule.....	343
hypervisor	
vGW Series kernel module.....	12
hypervisors	
overview.....	12

I

ICMPv6P.....	106
IDP Series interoperation.....	367
Image Enforcer	
Enforcer Profiles tab.....	172
Gold Image.....	171
Image Enforcer tab.....	171
installation	
prerequisites and resource requirements.....	15
Introspection module	
Applications feature.....	165
Image Enforcer.....	171, 172
Scan Status feature.....	177
Scheduling feature.....	176
VMs tab.....	168
IPv4	
defining network objects.....	316
IPv6	
address.....	49
communication with Juniper Networks IPv4	
servers.....	63
defining network objects.....	316
dual stack, IPv6 and IPv4.....	62
header fields.....	49
mixed component versions.....	60
overview.....	49
vGW Security Design VM modules.....	58

J

Junoscript XML interface.....360

K

kernel module.....12

L

licenses

evaluation.....208

features.....208

managing.....209

requirements.....208

understanding.....207

wizard.....209

log collection.....333

M

machines.....277

Main module.....77

manuals

comments on.....xxiv

multi-center feature.....244

N

Netflow

to Juniper STRM devices.....364

network connectivity.....15

Network module.....85

used with Firewall module.....90

network objects.....316

Network Settings323

Networks page.....316

NTP.....23

NTP services.....15

O

Open Virtualization Format (OVF)

OVA template.....23, 24

OVA bundled method.....24

downloading package.....25

OVA template.....23

OVA template method

single method for vGW Security Design

VM.....34

single OVA install method for vGW Security

VM.....36

overview.....3

P

policy

for multiple vNICs on the same VM.....258

Policy per vNIC

settings.....258

port groups.....15

Primary-level entry

secondary-level entry.....110, 247

Primary-level entry only.....110, 247

Protocols

ICPMv6.....106

protocols.....299

R

reports.....195

Reports module

AntiVirus.....202

resource requirements

DNS services.....15

network connectivity.....15

NTP services.....15

port groups.....15

vCenter.....15

virtual appliances.....15

virtual devices.....15

vNICs.....15

vSphere.....15

vSwitches.....15

Web browsers.....15

S

Security Settings

Global Settings.....282

Groups.....300

security zones

interfaces.....37

Settings module

Application Settings.....206

machines.....277

Network Settings323

overview.....203

settings module

multi-center.....244

Smart Groups.....305, 307

Groups.....300

SRX Series devices.....362, 363

SRX Series zones

SRX Series interoperability.....362

VM-to-zone mapping.....363

status	
appliance status.....	207
database status.....	207
status and status icons.....	77
STRM Series devices.....	364
support, technical See technical support	
Syslog	
to Juniper STRM devices.....	364
T	
taskbar.....	71
technical support	
contacting JTAC.....	xxiv
time synchronization.....	23
V	
vCenter.....	15
settings.....	209
vGW AntiVirus	
On-Access scanning.....	145
vGW Security Design	
modules.....	6
vGW Security Design VM	
Firewall module.....	90, 93
in mixed component version IPv6	
environment.....	60
Introspection module.....	176, 177
Applications feature.....	165
VMs tab.....	168
IPv6.....	58
Main module.....	77
navigation.....	69
Network module.....	85, 90
overview.....	67
predefined firewall policy rules	120
summary.....	3
taskbar.....	71
vGW Security VM Settings.....	286
vGW Security Design VM modules	
Compliance module.....	183
vGW Security VM	
in mixed component version IPv6	
environment.....	60
installation modes.....	22
overview.....	68
predefined firewall policy rules	120
single OVA install method.....	36
summary.....	3
vGW Security VM Settings.....	286
vGW Series.....	225, 227
VMware Auto Deploy support.....	225, 227
vGW Series and Juniper Networks interoperation	
IDP Series inter-operation.....	367
vGW Series interoperability	
creating SRX Series zone objects.....	362
Junoscript XML interface.....	360
populating zone address books with VM	
records.....	363
SRX Series devices.....	362
SRX Series zone objects.....	362, 363
STRM Series devices.....	364
virtual appliances.....	15
virtual devices	
sizes.....	15
VM-to-zone mapping.....	363
VMotion.....	13
VMsafe Firewall + Monitoring mode.....	22
VMsafe Monitoring mode.....	22
VMware.....	15
ESX/ESXi hosts.....	13
high availability.....	343
infrastructure and vGW.....	13
integrating vGW Series.....	209
preparation.....	21
VMotion.....	13
vSphere.....	13
vSphere operating system.....	13, 15
See also Virtual Center (vCenter)	
VMware VMsafe APIs.....	22
vNICs.....	15
vSphere.....	13
vSwitches.....	15
W	
Web browsers supported.....	15
Z	
zone address books.....	363
zones.....	362, 363

