paloalto TECHDOCS

VM-Series Deployment Guide

Version 10.0 (EoL)

docs.paloaltonetworks.com

Contact Information

Corporate Headquarters:
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com/company/contact-support

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal docs.paloaltonetworks.com.
- To search for a specific topic, go to our search page docs.paloaltonetworks.com/search.html.
- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Palo Alto Networks, Inc.
www.paloaltonetworks.com

© 2020-2022 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at www.paloaltonetworks.com/company/trademarks.html. All other marks mentioned herein may be trademarks of their respective companies.

Last Revised

July 16, 2022

Table of Contents

About the VM-Series Firewall 15

VM-Series Deployments 16
VM-Series in High Availability 18
Upgrade the VM-Series Firewall 20
Upgrade the PAN-OS Software Version (Standalone Version) 20
Upgrade the PAN-OS Software Version (HA Pair) 22
Upgrade the PAN-OS Software Version Using Panorama 26
Upgrade the PAN-OS Software Version (VM-Series for NSX) 29
Upgrade the VM-Series Model 35
Upgrade the VM-Series Model in an HA Pair 36
Downgrade a VM-Series Firewall to a Previous Release 37
VM-Series Plugin 39
Configure the VM-Series Plugin on the Firewall 39
Upgrade the VM-Series Plugin 40
Enable Jumbo Frames on the VM-Series Firewall 43
Hypervisor Assigned MAC Addresses 45
Custom PAN-OS Metrics Published for Monitoring 47
Interface Used for Accessing External Services on the VM-Series Firewall 49
PacketMMAP and DPDK Driver Support 50
Enable ZRAM on the VM-Series Firewall 52

License the VM-Series Firewall 53

VM-Series Firewall Licensing 54
License Types 54
Flexible vCPUs and Fixed Model Licensing 55
Flexible vCPUs and Fixed Model Deployment 57
Create a Support Account 59
Serial Number and CPU ID Format for the VM-Series Firewall 60
Licensing API 61
Manage the Licensing API Key 62
Use the Licensing API 62
Licensing API Error Codes 66
Install a License API Key 67
Use Panorama-Based Software Firewall License Management 68
Software NGFW Credits 74
Maximum Limits Based on Memory 76
Activate Credits 85
Create a Deployment Profile 86

 ©2024 Palo Alto Networks, Inc.

Manage a Deployment Profile....87
Register the VM-Series Firewall (Software NGFW Credits)....89
Provision Panorama....92
Migrate Panorama to a FW-Flex License....93
Transfer Credits....94
Renew Your Software NGFW Credit License....96
Deactivate License (Software NGFW Credits)....97
Create and Apply a Subscription-Only Auth Code....99
Migrate to a Flexible VM-Series License....102
VM-Series Models....114
VM-Series System Requirements....115
CPU Oversubscription....117
VM-50 Lite Mode....118
VM-Series Model License Types....118
Activate VM-Series Model Licenses....130
Register the VM-Series Firewall....136
Install a Device Certificate on the VM-Series Firewall....139
Switch Between the BYOL and the PAYG Licenses....141
Switch Between VM-Series Model Licenses....143
Deactivate License(s)....146
Renew VM-Series Firewall License Bundles....151
What Happens When Licenses Expire?....154
Licenses for Cloud Security Service Providers (CSSPs)....157
Get the Auth Codes for CSSP License Packages....157
Register the VM-Series Firewall with a CSSP Auth Code....158
Add End-Customer Information for a Registered VM-Series Firewall....159
Set Up a VM-Series Firewall on an ESXi Server....163
Supported Deployments on VMware vSphere Hypervisor (ESXi)....164
VM-Series on ESXi System Requirements and Limitations....165
VM-Series on ESXi System Requirements....165
VM-Series on ESXi System Limitations....166
Install a VM-Series firewall on VMware vSphere Hypervisor (ESXi)....167
Plan the Interfaces for the VM-Series for ESXi....167
Provision the VM-Series Firewall on an ESXi Server....168
Perform Initial Configuration on the VM-Series on ESXi....171
Add Additional Disk Space to the VM-Series Firewall....172
Use VMware Tools on the VM-Series Firewall on ESXi and vCloud Air....174
Use vMotion to Move the VM-Series Firewall Between Hosts....175
Use the VM-Series CLI to Swap the Management Interface on ESXi....176
VM Monitoring on vCenter....177

 ©2024 Palo Alto Networks, Inc.

About VM Monitoring on VMware vCenter 177
Install the Panorama Plugin for VMware vCenter 178
Configure the Panorama Plugin for VMware vCenter 179
Troubleshoot ESXi Deployments 182
Basic Troubleshooting 182
Installation Issues 182
Licensing Issues 184
Connectivity Issues 185
Performance Tuning of the VM-Series for ESXi 187
Install the NIC Driver on ESXi 187
Enable DPDK on ESXi 188
Enable SR-IOV on ESXi 189
Enable ESXi VLAN Access Mode with SR-IOV 189
Enable Multi-Queue Support for NICs on ESXi 190
VNF Tuning for Performance 191

Set Up the VM-Series Firewall on vCloud Air 205

About the VM-Series Firewall on vCloud Air 206
Deployments Supported on vCloud Air 207
Deploy the VM-Series Firewall on vCloud Air 208

Set Up the VM-Series Firewall on VMware NSX 217

Set Up the VM-Series Firewall on VMware NSX-V 218
VM-Series for Firewall NSX-V Overview 218
VM-Series Firewall for NSX-V Deployment Checklist 227
Install the VMware NSX Plugin 230
Register the VM-Series Firewall as a Service on the NSX-V Manager 231
Deploy the VM-Series Firewall 241
Create Security Groups and Steering Rules 249
Apply Security Policies to the VM-Series Firewall 255
Steer Traffic from Guests that are not Running VMware Tools 261
What is Multi-NSX Manager Support on the VM-Series for NSX-V? 262
Dynamically Quarantine Infected Guests 266
Migrate Operations-Centric Configuration to Security-Centric Configuration 270
Add a New Host to Your NSX-V Deployment 274
Use Case: Shared Compute Infrastructure and Shared Security Policies 275
Use Case: Shared Security Policies on Dedicated Compute Infrastructure 280
Dynamic Address Groups—Information Relay from NSX-V Manager to Panorama 287
Set Up the VM-Series Firewall on VMware NSX-T (North-South) 291

 ©2024 Palo Alto Networks, Inc.

Supported Deployments of the VM-Series Firewall on VMware NSX-T (North-South)....291
Components of the VM-Series Firewall on NSX-T (North-South)....292
Deploy the VM-Series Firewall on NSX-T (North-South)....293
Extend Security Policy from NSX-V to NSX-T....307
Set Up the VM-Series Firewall on NSX-T (East-West)....309
Components of the VM-Series Firewall on NSX-T (East-West)....309
VM-Series Firewall on NSX-T (East-West) Integration....310
Supported Deployments of the VM-Series Firewall on VMware NSX-T (East-West)....312
Deploy the VM-Series Using the Operations-Centric Workflow....313
Deploy the VM-Series Using the Security-Centric Workflow....331
Delete a Service Definition from Panorama....360
Migrate from VM-Series on NSX-T Operation to Security Centric Deployment....361
Extend Security Policy from NSX-V to NSX-T....367
Use In-Place Migration to Move Your VM-Series from NSX-V to NSX-T....368
Use Migration Coordinator to Move Your VM-Series from NSX-V to NSX-T....374

Set Up the VM-Series Firewall on AWS....381

About the VM-Series Firewall on AWS....382
AWS EC2 Instance Types....382
VM-Series Firewall on AWS GovCloud....382
VM-Series Firewall on AWS China....383
VM-Series Firewall on AWS Outposts....383
AWS Terminology....383
Management Interface Mapping for Use with Amazon ELB....386
Performance Tuning for the VM-Series Firewall on AWS....387
Deployments Supported on AWS....389
Deploy the VM-Series Firewall on AWS....392
Obtain the AMI....392
Planning Worksheet for the VM-Series in the AWS VPC....397
Launch the VM-Series Firewall on AWS....399
Launch the VM-Series Firewall on AWS Outpost....408
Create a Custom Amazon Machine Image (AMI)....416
Encrypt EBS Volume for the VM-Series Firewall on AWS....419
Use the VM-Series Firewall CLI to Swap the Management Interface....421
Enable CloudWatch Monitoring on the VM-Series Firewall....422
Panorama Orchestrated Deployments in AWS....427
Prepare for an Orchestrated AWS Deployment....429
Orchestrate a VM-Series Firewall Deployment in AWS....435

 ©2024 Palo Alto Networks, Inc.

View the Deployment Status....442
Traffic Flow and Configurations....444
VM-Series Integration with an AWS Gateway Load Balancer....452
Manual Integration of the VM-Series with a Gateway Load Balancer....454
VM-Series Auto Scaling Group with AWS Gateway Load Balancer....469
High Availability for VM-Series Firewall on AWS....483
Overview of HA on AWS....483
IAM Roles for HA....484
HA Links....488
Heartbeat Polling and Hello Messages....488
Device Priority and Preemption....488
HA Timers....489
Configure Active/Passive HA on AWS Using a Secondary IP....489
Configure Active/Passive HA on AWS Using Interface Move....495
Migrate Active/Passive HA on AWS....498
Use Case: Secure the EC2 Instances in the AWS Cloud....504
Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC....519
Use Case: VM-Series Firewalls as GlobalProtect Gateways on AWS....524
Components of the GlobalProtect Infrastructure....524
Deploy GlobalProtect Gateways on AWS....525
Resource Monitoring on AWS....527
AWS Resource Monitoring with the AWS Plugin on Panorama....527
Set Up the AWS Plugin for VM Monitoring on Panorama....527
Auto Scaling VM-Series Firewalls with the Amazon ELB Service....544
VM-Series Auto Scaling Templates for AWS Version 2.0....546
VM-Series Auto Scaling Templates for AWS Version 2.1....583
List of Attributes Monitored on the AWS VPC....612
IAM Permissions Required for Monitoring the AWS VPC....614

Set Up the VM-Series Firewall on KVM....615

VM-Series on KVM—Requirements and Prerequisites....616
Options for Attaching the VM-Series on the Network....617
Prerequisites for VM-Series on KVM....618
Supported Deployments on KVM....623
Secure Traffic on a Single Host....623
Secure Traffic Across Linux hosts....623
Install the VM-Series Firewall on KVM....625
Install the VM-Series Firewall Using Virt-Manager....625
Install the VM-Series Firewall Using an ISO....632
Use the VM-Series CLI to Swap the Management Interface on KVM....635

 ©2024 Palo Alto Networks, Inc.

Enable the Use of a SCSI Controller....636
Verify PCI-ID for Ordering of Network Interfaces on the VM-Series Firewall.... 638
Performance Tuning of the VM-Series for KVM....640
Install KVM and Open vSwitch on Ubuntu 16.04.1 LTS....640
Enable Open vSwitch on KVM....641
Integrate Open vSwitch with DPDK....641
Enable SR-IOV on KVM....646
Enable VLAN Access Mode with SR-IOV....647
Enable Multi-Queue Support for NICs on KVM....648
Isolate CPU Resources in a NUMA Node on KVM....648
Set Up the VM-Series Firewall on Hyper-V....651
Supported Deployments on Hyper-V....652
Secure Traffic on a Single Hyper-V Host....652
Secure Traffic Across Multiple Hyper-V Hosts....652
System Requirements on Hyper-V....654
Linux Integration Services....655
Install the VM-Series Firewall on Hyper-V....656
Before You Begin....656
Performance Tuning of the VM-Series Firewall on Hyper-V....657
Provision the VM-Series Firewall on a Hyper-V host with Hyper-V Manager....657
Provision the VM-Series Firewall on a Hyper-V host with PowerShell....658
Perform Initial Configuration on the VM-Series Firewall....660
Set up the VM-Series Firewall on Azure....663
About the VM-Series Firewall on Azure....664
Azure Networking and VM-Series Firewall....664
Azure Security Center Integration....665
VM-Series Firewall Templates on Azure....667
Minimum System Requirements for the VM-Series on Azure....668
Support for High Availability on VM-Series on Azure....669
VM-Series on Azure Service Principal Permissions....669
Deployments Supported on Azure....673
Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template)....674
Deploy the VM-Series Firewall from the Azure China Marketplace (Solution Template)....684
Panorama Orchestrated Deployments in Azure....690
Prepare for an Orchestrated Deployment....692
Orchestrate a VM-Series Firewall Deployment in Azure....696
Create a Custom VM-Series Image for Azure....705

 ©2024 Palo Alto Networks, Inc.

Use Azure Security Center Recommendations to Secure Your Workloads....707
Deploy a VM-Series Firewall Based on an Azure Security Center Recommendation....707
Connect an Existing VM-Series Firewall From Azure Security Center....708
Use Panorama to Forward Logs to Azure Security Center....710
Deploy the VM-Series Firewall on Azure Stack....713
Enable Azure Application Insights on the VM-Series Firewall....717
Monitoring on Azure....720
About Monitoring on Azure....720
Set Up the Azure Plugin for Monitoring on Panorama....720
Attributes Monitored Using the Panorama Plugin on Azure....728
Set up Active/Passive HA on Azure....731
Set up Active/Passive HA on Azure (North-South & East-West Traffic)....732
Set up Active/Passive HA on Azure (East-West Traffic Only)....749
Use the ARM Template to Deploy the VM-Series Firewall....765
Deploy the VM-Series and Azure Application Gateway Template....769
VM-Series and Azure Application Gateway Template....770
Start Using the VM-Series & Azure Application Gateway Template....771
Secure Kubernetes Services on Azure....779
How Does the Panorama Plugin for Azure Secure Kubernetes Services?....779
Secure an AKS Cluster....783

Set Up the VM-Series Firewall on OpenStack....797

VM-Series Deployments in OpenStack....798
Basic Gateway....798
Service Chaining and Service Scaling....798
Components of the VM-Series for OpenStack Solution....800
Heat Template for a Basic Gateway Deployment....802
Heat Templates for Service Chaining and Service Scaling....805
Virtual Network....806
Virtual Machine....806
Service Template....807
Service Instance....808
IPAM....808
Service Policy....809
Alarm....809
Install the VM-Series Firewall in a Basic Gateway Deployment....811
Install the VM-Series Firewall with Service Chaining or Scaling....814

Set Up the VM-Series Firewall on Google Cloud Platform....819

About the VM-Series Firewall on Google Cloud Platform....820

 ©2024 Palo Alto Networks, Inc.

Google Cloud Platform and the VM-Series Firewall....820
Minimum System Requirements for the VM-Series Firewall....820
Supported Deployments on Google Cloud Platform....822
Internet Gateway....822
Segmentation Gateway....822
Hybrid IPSec VPN....823
Prepare to Set Up VM-Series Firewalls on Google Public Cloud....824
General Requirements....824
Install the VM-Series Plugin on Panorama....825
Install the Panorama Plugin for GCP....825
Prepare to Deploy from the GCP Marketplace....826
Create a Custom VM-Series Firewall Image for Google Cloud Platform....834
Deploy the VM-Series Firewall on Google Cloud Platform....837
Deploy the VM-Series Firewall from Google Cloud Platform Marketplace....837
Management Interface Swap for Google Cloud Platform Load Balancing....841
Use the VM-Series Firewall CLI to Swap the Management Interface....843
Enable Google Stackdriver Monitoring on the VM Series Firewall....844
Enable VM Monitoring to Track VM Changes on Google Cloud Platform (GCP)....846
Use Dynamic Address Groups to Secure Instances Within the VPC....849
Use Custom Templates or the gcloud CLI to Deploy the VM-Series Firewall....851
VM Monitoring with the Panorama Plugin for GCP....853
Configure VM Monitoring with the Panorama Plugin for GCP....853
Auto Scaling the VM-Series Firewall on Google Cloud Platform....860
Auto Scaling Components for Google Cloud Platform....860
Deploy GCP Auto Scaling Templates....872

Set Up a VM-Series Firewall on a Cisco ENCS Network....901

Plan Your Cisco ENCS Deployment....902
Prepare the VM-Series Firewall Image for Cisco ENCS....904
Convert a qcow2 File from the Graphical User Interface....904
Convert a qcow2 File from the Command Line Interface....906
Deploy the VM-Series Firewall on Cisco ENCS....911

Set up the VM-Series Firewall on Oracle Cloud Infrastructure....915

OCI Shape Types....916
Deployments Supported on OCI....917
Prepare to Set Up the VM-Series Firewall on OCI....918
Deploy the VM-Series Firewall From the Oracle Cloud Marketplace....921
Configure Active/Passive HA on OCI....926

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on Alibaba Cloud.....935

VM-Series Firewall on Alibaba Cloud.....936
Minimum System Requirements for the VM-Series Firewall on Alibaba Cloud.....937
VM-Series Firewall Software Requirements.....937
Alibaba Cloud Instance Type Recommendations for the VM-Series Firewall.....937
Alibaba Cloud CLI.....937
Prepare to Deploy the VM-Series Firewall on Alibaba Cloud.....938
Choose Licenses and Plan Networks.....938
Prepare to Use the Aliyun Command Line Interface.....939
Deploy the VM-Series Firewall on Alibaba Cloud.....940
Create a VPC and Configure Networks.....940
Create and Configure the VM-Series Firewall.....943
Secure North-South Traffic on Alibaba Cloud.....947
Configure Load Balancing on Alibaba Cloud.....951

Set Up a Firewall in Cisco ACI.....955

Palo Alto Networks Firewall Integration with Cisco ACI.....956
Service Graph Templates.....957
Multi-Context Deployments.....957
Prepare Your ACI Environment for Integration.....958
Integrate the Firewall with Cisco ACI in Network Policy Mode.....959
Deploy the Firewall to Secure East-West Traffic in Network Policy Mode.....959
Deploy the Firewall to Secure North-South Traffic in Network Policy Mode.....974
Endpoint Monitoring in Cisco ACI.....991
Install the Panorama Plugin for Cisco ACI.....992
Configure the Cisco ACI Plugin.....993
Panorama Plugin for Cisco ACI Dashboard.....997

Set Up the VM-Series Firewall on Cisco CSP.....1001

VM-Series on Cisco CSP System Requirements.....1002
Deploy the VM-Series Firewall on Cisco CSP.....1003

Endpoint Monitoring for Cisco TrustSec.....1005

Panorama Plugin for Cisco TrustSec.....1006
Bulk Sync.....1006
PubSub.....1007
Install the Panorama Plugin for Cisco TrustSec.....1008
Configure the Panorama Plugin for Cisco TrustSec.....1009

 ©2024 Palo Alto Networks, Inc.

Troubleshoot the Panorama Plugin for Cisco TrustSec ... 1016
Plugin Status Commands ... 1016
Debug Commands ... 1016
Debug Logs ... 1017

Set Up the VM-Series Firewall on Nutanix AHV ... 1019

VM Monitoring on Nutanix ... 1020
About VM Monitoring on Nutanix ... 1020
Install the Panorama Plugin for Nutanix ... 1021
Configure the Panorama Plugin for Nutanix ... 1022

Bootstrap the VM-Series Firewall ... 1027

Choose a Bootstrap Method ... 1028
Basic Configuration ... 1029
Complete Configuration ... 1033
VM-Series Firewall Bootstrap Workflow ... 1034
Bootstrap Package ... 1036
Bootstrap Package Structure ... 1036
Bootstrap Package Delivery ... 1036
Bootstrap Configuration Files ... 1038
init-cfg.txt ... 1038
bootstrap.xml ... 1038
Generate the VM Auth Key on Panorama ... 1039
Create the init-cfg.txt File ... 1041
init-cfg.txt File Components ... 1042
Sample init-cfg.txt File ... 1045
Create the bootstrap.xml File ... 1047
Prepare the Licenses for Bootstrapping ... 1048
Prepare the Bootstrap Package ... 1049
Bootstrap the VM-Series Firewall on AWS ... 1051
Bootstrap the VM-Series Firewall on Azure ... 1055
Bootstrap the VM-Series Firewall on ESXi ... 1059
Bootstrap the VM-Series Firewall on ESXi with an ISO ... 1059
Bootstrap the VM-Series Firewall on ESXi with a Block Storage Device ... 1059
Bootstrap the VM-Series Firewall on Google Cloud Platform ... 1061
Bootstrap the VM-Series Firewall on Hyper-V ... 1063
Bootstrap the VM-Series Firewall on Hyper-V with an ISO ... 1063
Bootstrap the VM-Series Firewall on Hyper-V with a Block Storage Device ... 1063
Bootstrap the VM-Series Firewall on KVM ... 1066
Bootstrap the VM-Series Firewall on KVM with an ISO ... 1066

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on KVM With a Block Storage Device.. 1066
Verify Bootstrap Completion..... 1068
Bootstrap Errors..... 1069

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

About the VM-Series Firewall

The Palo Alto Networks VM-Series firewall is the virtualized form of the Palo Alto Networks next-generation firewall. It is positioned for use in a virtualized or cloud environment where it can protect and secure east-west and north-south traffic.

- VM-Series Deployments
- VM-Series in High Availability
- Upgrade the VM-Series Firewall
- VM-Series Plugin
- Enable Jumbo Frames on the VM-Series Firewall
- Hypervisor Assigned MAC Addresses
- Custom PAN-OS Metrics Published for Monitoring
- Interface Used for Accessing External Services on the VM-Series Firewall
- PacketMMAP and DPDK Driver Support
- Enable ZRAM on the VM-Series Firewall

VM-Series Deployments

The VM-Series firewall can be deployed on the following platforms:

- **VM-Series for VMware vSphere Hypervisor (ESXi) and vCloud Air**

  You can deploy any VM-Series model as a guest virtual machine on VMware ESXi; ideal for cloud or networks where virtual form factor is required.

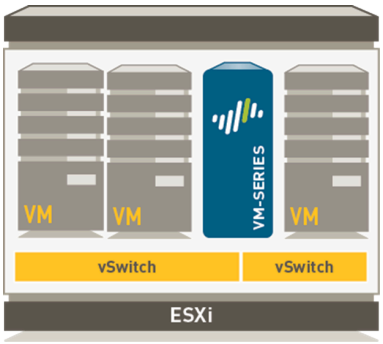

  For details, see Set Up a VM-Series Firewall on an ESXi Server and Set Up the VM-Series Firewall on vCloud Air.

- **VM-Series on VMware NSX-V**

  The VM-100, VM-200, VM-300, VM-500, or VM-1000-HV is deployed as a network introspection service with VMware NSX, and Panorama. This deployment is ideal for east-west traffic inspection, and it also can secure north-south traffic.

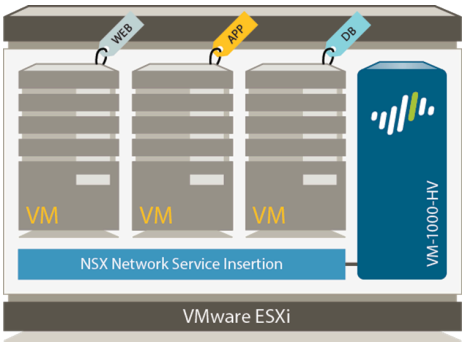

  For details, see Set Up the VM-Series Firewall on VMware NSX-V.

- **VM-Series on VMware NSX-T**

  You can deploy the VM-100, VM-300, VM-500, or VM-700 in your NSX-T environment.

  For details, see Set Up the VM-Series Firewall on VMware NSX-T (North-South).

- **VM-Series for Amazon Web Services (AWS)**

  You can deploy any VM-Series model, except the VM-50, on EC2 instances on the AWS Cloud.

  For details, see Set Up the VM-Series Firewall on AWS.

 ©2024 Palo Alto Networks, Inc.

- **VM-Series for Google Cloud Platform**

  You can deploy any VM-Series model, except the VM-50 and the VM-50 Lite on Google Compute Engine instances. For details, see Set Up the VM-Series Firewall on Google Cloud Platform.

- **VM-Series for Kernel Virtualization Module (KVM)**

  You can deploy any VM-Series model on a Linux server that is running the KVM hypervisor. For details, see Set Up the VM-Series Firewall on KVM.

- **VM-Series for Microsoft Hyper-V**

  You can deploy any VM-Series model on a Windows Server 2012 R2 server with the Hyper-V role add-on enabled or a standalone Hyper-V 2012 R2 server. For details, see Set Up the VM-Series Firewall on Hyper-V.

- **VM-Series for Microsoft Azure**

  You can deploy any VM-Series model, except the VM-50, on the Azure VNet.

  For details, see Set up the VM-Series Firewall on Azure.

 ©2024 Palo Alto Networks, Inc.

VM-Series in High Availability

High availability (HA) is a configuration in which two firewalls are placed in a group and their configuration is synchronized to prevent a single point of failure on your network. A heartbeat connection between the firewall peers ensures seamless failover in the event that a peer goes down. Setting up the firewalls in a two-device cluster provides redundancy and allows you to ensure business continuity. In an HA configuration on the VM-Series firewalls, both peers must be deployed on the same type of hypervisor, have identical hardware resources (such as CPU cores/network interfaces) assigned to them, and have the set same of licenses/subscriptions. For general information about HA on Palo Alto Networks firewalls, see High Availability.

The VM-Series firewalls support stateful active/passive or active/active high availability with session and configuration synchronization. The active/active deployment is supported in virtual wire and Layer 3 deployments on some private cloud hypervisors, and is recommended only if each firewall needs its own routing instances and you require full, real-time redundancy out of both firewalls all the time. To configure the VM-Series firewall as an HA pair, see Configure Active/Passive HA and Configure Active/Active HA.

If you are deploying the VM-Series firewall in the public cloud, such as on the Amazon Web Services (AWS) or Azure, you can use the traditional active/passive HA configuration; see High Availability for VM-Series Firewall on AWS and Set up Active/Passive HA on Azure. Alternatively, because of the innate differences in how resource or region redundancy is built into the cloud infrastructure as compared to a private data center, to take advantage of native cloud services and build a resilient architecture that maximizes uptime, see

- AWS— Auto Scale VM-Series Firewalls with the Amazon ELB to deploy multiple firewalls across two or more Availability Zones within a VPC.
- Azure— VM-Series and Azure Application Gateway Template Parameters.

Features/ Links Supported	ESX	KVM	AWS	NSX V	NSX T (N/S)	Hyper-V	Azure	GCP	OCI
Active/Passive HA	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes
Active/Active HA	Yes	Yes	No	No	No	Yes	No	No	No
HA 1	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
HA2—(session synchronization and keepalive)	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
HA3	Yes	Yes	No	No	No	Yes	No	No	No

HA1 and HA2 support for the VM-Series on GCP requires PAN-OS 10.0x or later and VM-Series plugin 2.0.5 or later.

 ©2024 Palo Alto Networks, Inc.

High availability for the VM-Series firewall on NSX-T (E/W) is achieved through the NSX-T feature called service health check. This NSX-T feature allows you to simulate high availability in the case of a service instance failing. When configured with the VM-Series firewall, if a VM-Series service instance fails, any traffic directed to that firewall is redirect to another firewall instance in the cluster (for service cluster deployments) or a firewall instance on another host (for host-based deployments). See Configure the Service Definition on Panorama for the VM-Series firewall on NSX-T (E/W) for more information.

 ©2024 Palo Alto Networks, Inc.

Upgrade the VM-Series Firewall

Upgrading the PAN-OS version or VM-Series model allows you to add the latest features and fixes that help improve the security capabilities and performance of your firewalls.

The standard PAN-OS release is just that; the normal version of PAN-OS that can be installed on all Palo Alto Networks firewalls. The PAN-OS XFR releases are for VM-Series firewalls only and can include new features and bug fixes for VM-Series firewalls. If you install a PAN-OS XFR image on the VM-Series firewalls, the features and fixes are not available in PAN-OS versions that are earlier than the software version you have installed.

Because XFR images include features and fixes that are specific to VM-Series firewalls, if you upgrade to an XFR release, you must stay on XFR releases to keep XFR specific features until the next major PAN-OS release; all the fixes and capabilities available in XFR will be cumulatively rolled into the next major PAN-OS release.

- Upgrade the PAN-OS Software Version (Standalone Version)
- Upgrade the PAN-OS Software Version (HA Pair)
- Upgrade the PAN-OS Software Version Using Panorama
- Upgrade the PAN-OS Software Version (VM-Series for NSX)
- Upgrade the VM-Series Model
- Upgrade the VM-Series Model in an HA Pair
- Downgrade a VM-Series Firewall to a Previous Release

For instructions on installing your VM-Series firewall, see VM-Series Deployments.

Verify the VM-Series System Requirements for your firewall model before you upgrade. If your firewall has less than 5.5GB memory, the system capacity (number of sessions, rules, security zones, address objects, etc) on the firewall will be limited to that of the VM-50 Lite.

Upgrade the PAN-OS Software Version (Standalone Version)

Review the new features, addressed issues, and known issues and then use the following procedure to upgrade a firewall that is not in an HA configuration.

To avoid impacting traffic, plan to upgrade within the outage window. Ensure the firewall is connected to a reliable power source. A loss of power during an upgrade can make the firewall unusable.

STEP 1 | Verify that enough hardware resources are available to the VM-Series firewall.

Refer to the VM-Series System Requirements to see the resource requirements for each VM-Series model. Allocate additional hardware resources before continuing the upgrade process; the process for assigning additional hardware resources differs on each hypervisor.

If the VM-Series firewall does not have the required resources for the model, it defaults to the capacity associated with the VM-50.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | From the web interface, navigate to **Device** > **Licenses** and make sure you have the correct VM-Series firewall license and that the license is activated.

On the VM-Series firewall standalone version, navigate to **Device** > **Support** and make sure that you have activated the support license.

STEP 3 | Save a backup of the current configuration file.

Although the firewall automatically creates a configuration backup, it is a best practice to create and externally store a backup before you upgrade.

1. Select **Device** > **Setup** > **Operations** and click **Export named configuration snapshot**.
2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.
3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 4 | If you have enabled User-ID, after you upgrade, the firewall clears the current IP address-to-username and group mappings so that they can be repopulated with the attributes from the User-ID sources. To estimate the time required for your environment to repopulate the mappings, run the following CLI commands on the firewall.

- For IP address-to-username mappings:
  - **`show user user-id-agent state all`**
  - **`show user server-monitor state all`**
- For group mappings: **`show user group-mapping statistics`**

STEP 5 | Ensure that the firewall is running the latest content release version.

1. Select **Device** > **Dynamic Updates** and see which **Applications** or **Applications and Threats** content release version is Currently Installed.
2. If the firewall is not running the minimum required content release version or a later version required for PAN-OS, **Check Now** to retrieve a list of available updates.
3. Locate and **Download** the desired content release version.

   After you successfully download a content update file, the link in the Action column changes from **Download** to **Install** for that content release version.
4. **Install** the update.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Upgrade the VM-Series plugin.

1. Before upgrading, check the latest Release Notes for details on whether a new VM-Series plugin affects your environment.

   For example, suppose a new VM-Series plugin version only includes AWS features. To take advantage of the new features, you must update the plugin on your VM-Series firewall instances on AWS.

   Do not install an upgrade that does not apply to your environment.

2. Log in to the VM-Series firewall and check the dashboard to view the plugin version.
3. Select **Device** > **Plugins** to view the plugin version. Use **Check Now** to check for updates.
4. Select the version of the plugin and click **Install** in the Action column to install the plugin.

STEP 7 | Upgrade PAN-OS.

*If your firewall does not have internet access from the management port, you can download the software image from the Palo Alto Networks Customer Support Portal and then manually **Upload** it to your firewall.*

1. Select **Device** > **Software** and click **Check Now** to display the latest PAN-OS updates.
2. Locate and **Download** the target PAN-OS version.
3. After you download the image (or, for a manual upgrade, after you upload the image), **Install** the image.
4. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and click **Reboot Device**.

   At this point, the firewall clears the User-ID mappings, then connects to the User-ID sources to repopulate the mappings.
5. If you have enabled User-ID, use the following CLI commands to verify that the firewall has repopulated the IP address-to-username and group mappings before allowing traffic.
   - `show user ip-user-mapping all`
   - `show user group list`
6. If you are upgrading to an XFR release for the first time, repeat this step to upgrade to the corresponding XFR release.

STEP 8 | Verify that the firewall is passing traffic.

Select **Monitor** > **Session Browser** and verify that you are seeing new sessions.

Upgrade the PAN-OS Software Version (HA Pair)

Use the following procedure to upgrade a pair of firewalls in a high availability (HA) configuration. This procedure applies to both active/passive and active/active configurations.

 ©2024 Palo Alto Networks, Inc.

To avoid downtime when upgrading firewalls that are in a high availability (HA) configuration, update one HA peer at a time: For active/active firewalls, it doesn't matter which peer you upgrade first (though for simplicity, this procedure shows you how to upgrade the active-secondary peer first). For active/passive firewalls, you must upgrade the passive peer first, suspend the active peer (fail over), update the active peer, and then return that peer to a functional state (fail back). To prevent failover during the upgrade of the HA peers, you must make sure preemption is disabled before proceeding with the upgrade. You only need to disable preemption on one peer in the pair.

To avoid impacting traffic, plan to upgrade within the outage window. Ensure the firewalls are connected to a reliable power source. A loss of power during an upgrade can make firewalls unusable.

STEP 1 | Verify that enough hardware resources are available to the VM-Series firewall.

Refer to the VM-Series System Requirements to see the resource requirements for each VM-Series model. Allocate additional hardware resources before continuing the upgrade process; the process for assigning additional hardware resources differs on each hypervisor.

If the VM-Series firewall does not have the required resources for the model, it defaults to the capacity associated with the VM-50.

STEP 2 | From the web interface, navigate to **Device** > **Licenses** and make sure you have the correct VM-Series firewall license and that the license is activated.

On the VM-Series firewall standalone version, navigate to **Device** > **Support** and make sure that you have activated the support license.

STEP 3 | Save a backup of the current configuration file.

Although the firewall automatically creates a backup of the configuration, it is a best practice to create and externally store a backup before you upgrade.

Perform these steps on each firewall in the pair:

1. Select **Device** > **Setup** > **Operations** and click **Export named configuration snapshot**.
2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.
3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 4 | If you have enabled User-ID, after you upgrade, the firewall clears the current IP address-to-username and group mappings so that they can be repopulated with the attributes from the User-ID sources. To estimate the time required for your environment to repopulate the mappings, run the following CLI commands on the firewall.

- For IP address-to-username mappings:
  - **`show user user-id-agent state all`**
  - **`show user server-monitor state all`**
- For group mappings: **`show user group-mapping statistics`**

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Ensure that each firewall in the HA pair is running the latest content release version.

Refer to the release notes for the minimum content release version you must install for a PAN-OS 9.1 release. Make sure to follow the Best Practices for Application and Threat Updates.

1. Select **Device** > **Dynamic Updates** and check which **Applications** or **Applications and Threats** to determine which update is Currently Installed.
2. If the firewalls are not running the minimum required content release version or a later version required for the software version you are installing, **Check Now** to retrieve a list of available updates.
3. Locate and **Download** the desired content release version.

   After you successfully download a content update file, the link in the Action column changes from **Download** to **Install** for that content release version.
4. **Install** the update. You must install the update on both peers.

STEP 6 | Upgrade the VM-Series plugin.

1. Before upgrading, check the latest Release Notes for details on whether a new VM-Series plugin affects your environment.

   For example, suppose a new VM-Series plugin version only includes AWS features. To take advantage of the new features, you must update the plugin on your VM-Series firewall instances on AWS.

   Do not install an upgrade that does not apply to your environment.
2. Log in to the VM-Series firewall and check the dashboard to view the plugin version.
3. Select **Device** > **Plugins** to view the plugin version. Use **Check Now** to check for updates.
4. Select the version of the plugin and click **Install** in the Action column to install the plugin.

   When installing the plugin on VM-Series firewalls in an HA pair, install the higher version VM-Series plugin on the active peer before the passive peer. After installing the plugin on the active peer it transitions the passive peer to a non-functional state. Installing the plugin on the passive peer returns the passive peer to a functional state.

STEP 7 | Disable preemption on the first peer in each pair. You only need to disable this setting on one firewall in the HA pair but ensure that the commit is successful before you proceed with the upgrade.

1. Select **Device** > **High Availability** and edit the **Election Settings**.
2. If enabled, disable (clear) the **Preemptive** setting and click **OK**.
3. **Commit** the change.

STEP 8 | Install the PAN-OS release on the first peer. If you are upgrading to an XFR release, install the version that corresponds to the XFR release.

To minimize downtime in an active/passive configuration, upgrade the passive peer first. For an active/active configuration, upgrade the secondary peer first. As a best practice, if you are

 ©2024 Palo Alto Networks, Inc.

using an active/active configuration, we recommend upgrading both peers during the same maintenance window.

If you want to test that HA is functioning properly before the upgrade, consider upgrading the active peer in an active/passive configuration first to ensure that failover occurs without incident.

1. On the first peer, select **Device** > **Software** and click **Check Now** for the latest updates.
2. Locate and **Download** the target PAN-OS version.

   *If your firewall does not have internet access from the management port, you can download the software image from the Palo Alto Networks Support Portal and then manually **Upload** it to your firewall.*
3. After you download the image (or, for a manual upgrade, after you upload the image), **Install** the image.
4. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
5. After the device finishes rebooting, view the High Availability widget on the **Dashboard** and verify that the device you just upgraded is still the passive or active-secondary peer in the HA configuration.

STEP 9 | Install the PAN-OS release on the second peer. If you are upgrading to an XFR release, install the version that corresponds to the XFR release.

1. (Active/passive configurations only) Suspend the active peer so that HA fails over to the peer you just upgraded.
   1. On the active peer, select **Device** > **High Availability** > **Operational Commands** and click **Suspend local device**.
   2. View the High Availability widget on the **Dashboard** and verify that the state changes to **Passive**.
   3. On the other peer, verify that it is active and is passing traffic (**Monitor** > **Session Browser**).
2. On the second peer, select **Device** > **Software** and click **Check Now** for the latest updates.
3. Locate and **Download** the target PAN-OS version.
4. After you download the image, **Install** it.
5. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
6. (Active/passive configurations only) From the CLI of the peer you just upgraded, run the following command to make the firewall functional again:

   ```
   request high-availability state functional
   ```

 ©2024 Palo Alto Networks, Inc.

STEP 10 | (PAN-OS XFR upgrade only) Upgrade the first peer and second peer to PAN-OS XFR by repeating Step 8 and Step 9.

STEP 11 | Verify that both peers are passing traffic as expected.

In an active/passive configuration, only the active peer should be passing traffic; both peers should be passing traffic in an active/active configuration.

Run the following CLI commands to confirm that the upgrade succeeded:

- (Active peers only) To verify that active peers are passing traffic, run the **`show session all`** command.
- To verify session synchronization, run the **`show high-availability interface ha2`** command and make sure that the Hardware Interface counters on the CPU table are increasing as follows:
  - In an active/passive configuration, only the active peer shows packets transmitted; the passive peer will show only packets received.

    If you enabled HA2 keep-alive, the hardware interface counters on the passive peer will show both transmit and receive packets. This occurs because HA2 keep-alive is bi-directional, which means that both peers transmit HA2 keep-alive packets.
  - In an active/active configuration, you will see packets received and packets transmitted on both peers.

STEP 12 | If you disabled preemption prior to the upgrade, re-enable it now.

1. Select **Device** > **High Availability** and edit the **Election Settings**.
2. Select **Preemptive** and click **OK**.
3. **Commit** the change.

Upgrade the PAN-OS Software Version Using Panorama

Use the following procedure to upgrade firewalls that you manage with Panorama. This procedure applies to standalone firewalls and firewalls deployed in a high availability (HA) configuration.

If Panorama is unable to connect directly to the update server, follow the procedure for deploying updates to firewalls when Panorama is not internet-connected so that you can manually download images to Panorama and then distribute the images to firewalls.

Before you can upgrade firewalls from Panorama, you must:

- Make sure Panorama is running the same or a later PAN-OS version than you are upgrading to. You must upgrade Panorama and its Log Collectors to 9.1 before upgrading the managed firewalls to this version. In addition, when upgrading Log Collectors to 9.1, you must upgrade all Log Collectors at the same time due to changes in the logging infrastructure.
- Plan for an extended maintenance window of up to six hours when upgrading Panorama to 9.1. This release includes significant infrastructure changes, which means that the Panorama upgrade will take longer than in previous releases.

 ©2024 Palo Alto Networks, Inc.

- ❑ Ensure that firewalls are connected to a reliable power source. A loss of power during an upgrade can make a firewall unusable.

STEP 1 | After upgrading Panorama, commit and push the configuration to the firewalls you are planning to upgrade.

STEP 2 | Verify that enough hardware resources are available to the VM-Series firewall.

Refer to the VM-Series System Requirements to see the resource requirements for each VM-Series model. Allocate additional hardware resources before continuing the upgrade process; the process for assigning additional hardware resources differs on each hypervisor.

If the VM-Series firewall does not have the required resources for the model, it defaults to the capacity associated with the VM-50.

STEP 3 | From the web interface, navigate to **Device** > **Licenses** and make sure you have the correct VM-Series firewall license and that the license is activated.

On the VM-Series firewall standalone version, navigate to **Device** > **Support** and make sure that you have activated the support license.

STEP 4 | Save a backup of the current configuration file on each managed firewall you plan to upgrade.

Although the firewall automatically creates a configuration backup, it is a best practice to create and externally store a backup before you upgrade.

1. From the Panorama web interface, select **Panorama** > **Setup** > **Operations** and click **Export Panorama and devices config bundle** to generate and export the latest configuration backup of Panorama and of each managed appliance.
2. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 5 | Update the content release version on the firewalls you plan to upgrade.

Refer to the Release Notes for the minimum content release version required for PAN-OS 9.1. Make sure to follow the Best Practices for Application and Threat Updates when deploying content updates to Panorama and managed firewalls.

1. Select **Panorama** > **Device Deployment** > **Dynamic Updates** and **Check Now** for the latest updates. If an update is available, the Action column displays a **Download** link.
2. If not already installed, **Download** the latest content release version.
3. Click **Install**, select the firewalls on which you want to install the update, and click **OK**. If you are upgrading HA firewalls, you must update content on both peers.

STEP 6 | (HA firewall upgrades only) If you will be upgrading firewalls that are part of an HA pair, disable preemption. You need only disable this setting on one firewall in each HA pair.

1. Select **Device** > **High Availability** and edit the **Election Settings**.
2. If enabled, disable (clear) the **Preemptive** setting and click **OK**.
3. **Commit** your change. Make sure the commit is successful before you proceed with the upgrade.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Download the target PAN-OS release image.

1. Select **Panorama** > **Device Deployment** > **Software** and **Check Now** for the latest release versions.
2. **Download** the firewall-specific file (or files) for the release version to which you are upgrading. You must download a separate installation file for each firewall model (or firewall series) that you intend to upgrade.

STEP 8 | Install the PAN-OS software update on the firewalls.

1. Click **Install** in the Action column that corresponds to the firewall models you want to upgrade.
2. In the Deploy Software file dialog, select all firewalls that you want to upgrade. To reduce downtime, select only one peer in each HA pair. For active/passive pairs, select the passive peer; for active/active pairs, select the active-secondary peer.
3. (HA firewall upgrades only) Make sure **Group HA Peers** is not selected.
4. Select **Reboot device after install**.
5. To begin the upgrade, click **OK**.
6. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
7. After the firewalls finish rebooting, select **Panorama** > **Managed Devices** and verify the Software Version is 9.1.0 for the firewalls you upgraded. Also verify that the HA status of any passive firewalls you upgraded is still passive.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | (HA firewall upgrades only) Upgrade the second HA peer in each HA pair.

1. (Active/passive upgrades only) Suspend the active device in each active/passive pair you are upgrading.
   1. Switch context to the active firewall.
   2. In the High Availability widget on the **Dashboard**, verify that **Local** firewall state is **Active** and the **Peer** is **Passive**).
   3. Select **Device** > **High Availability** > **Operational Commands** > **Suspend local device**.
   4. Go back to the High Availability widget on the **Dashboard** and verify that **Local** changed to **Passive** and **Peer** changed to **Active**.
2. Go back to the Panorama context and select **Panorama** > **Device Deployment** > **Software**.
3. Click **Install** in the Action column that corresponds to the firewall models of the HA pairs you are upgrading.
4. In the Deploy Software file dialog, select all firewalls that you want to upgrade. This time, select only the peers of the HA firewalls you just upgraded.
5. Make sure **Group HA Peers** is not selected.
6. Select **Reboot device after install**.
7. To begin the upgrade, click **OK**.
8. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, select **Device** > **Setup** > **Operations** and **Reboot Device**.
9. (Active/passive upgrades only) From the CLI of the peer you just upgraded, run the following command to make the firewall functional again:

   ```
   request high-availability state functional
   ```

STEP 10 | (PAN-OS XFR upgrade only) Upgrade the first peer and second peer to PAN-OS XFR by repeating Step 8 and Step 9.

STEP 11 | Verify the software and content release version running on each managed firewall.

1. On Panorama, select **Panorama** > **Managed Devices**.
2. Locate the firewalls and review the content and software versions in the table.

   For HA firewalls, you can also verify that the HA Status of each peer is as expected.

STEP 12 | (HA firewall upgrades only) If you disabled preemption on one of your HA firewalls before you upgraded, then edit the **Election Settings** (**Device** > **High Availability**) and re-enable the **Preemptive** setting for that firewall and then **Commit** the change.

Upgrade the PAN-OS Software Version (VM-Series for NSX)

Choose the upgrade method that best suits your deployment.

 ©2024 Palo Alto Networks, Inc.

- Upgrade the VM-Series for NSX During a Maintenance Window—use this option to upgrade the VM-Series firewall during a maintenance window without changing the OVF URL in the service definition.
- Upgrade the VM-Series for NSX Without Disrupting Traffic—use this option to upgrade the VM-Series firewall without disrupting service to the guest VMs or changing the OVF URL in the service definition.

The following graphics displays the currently supported combinations of Panorama and the Panorama plugin for VMware NSX, as well as the upgrade paths you are required to follow to upgrade successfully.

- Each box below represents a supported combination.
- When upgrading the Panorama plugin for NSX or Panorama in an HA pair, upgrade the passive Panorama peer first, followed by the active HA peer.

Before upgrading your VM-Series for VMware NSX deployment, review the upgrade paths shown below to understand the upgrade steps to arrive at the plugin and PAN-OS combination that best suits your environment.

 ©2024 Palo Alto Networks, Inc.

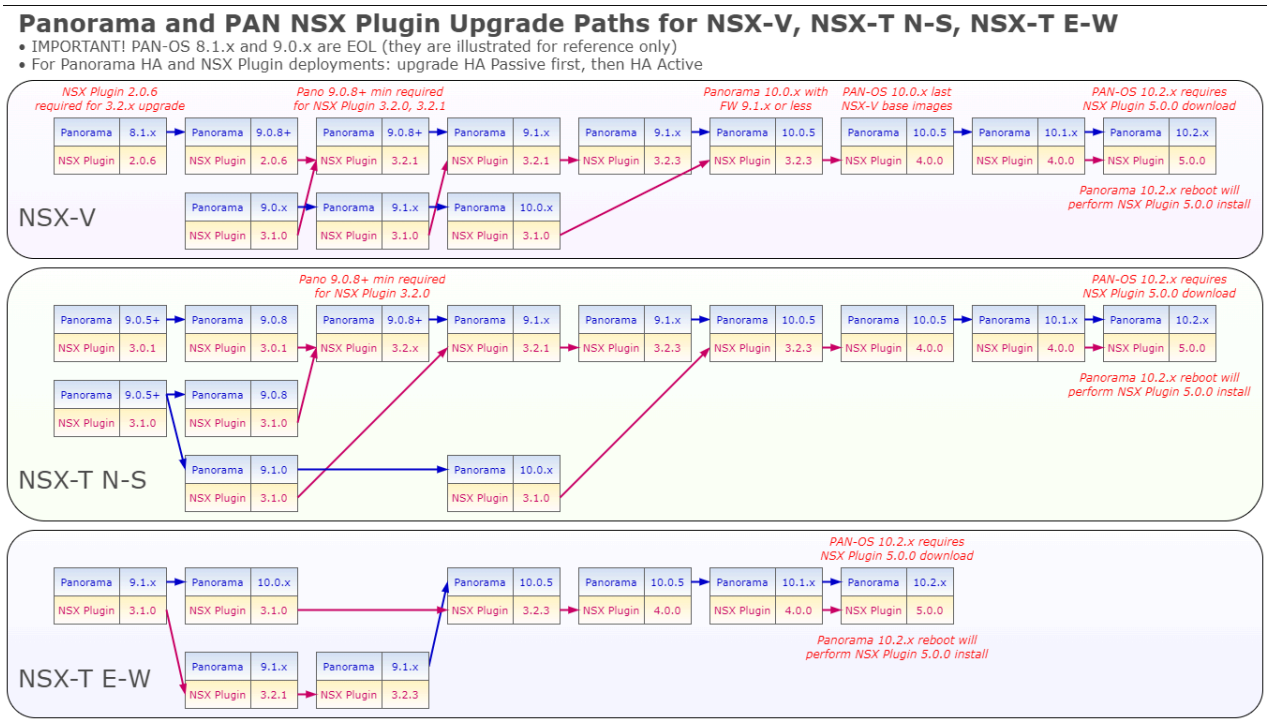

Upgrade the VM-Series for NSX During a Maintenance Window

For the VM-Series Firewall NSX edition, use Panorama to upgrade the software version on the firewalls.

STEP 1 | Review the VM-Series for VMware NSX upgrade paths.

STEP 2 | Allocate additional hardware resources to your VM-Series firewall.

Verify that enough hardware resources are available to the VM-Series firewall. Refer to the VM-Series System Requirements to see the new resource requirements for each VM-Series model. Allocate additional hardware resources before continuing the upgrade process. The process for assigning additional hardware resources differs on each hypervisor.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Save a backup of the current configuration file on each managed firewall that you plan to upgrade.

Although the firewall will automatically create a backup of the configuration, it is a best practice to create a backup prior to upgrade and store it externally.

1. Select **Device** > **Setup** > **Operations** and click **Export Panorama and devices config bundle**. This option is used to manually generate and export the latest version of the configuration backup of Panorama and of each managed device.
2. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 4 | Check the Release Notes to verify the Content Release version required for the PAN-OS version.

The firewalls you plan to upgrade must be running the Content Release version required for the PAN-OS version.

1. Select **Panorama** > **Device Deployment** > **Dynamic Updates**.
2. Check for the latest updates. Click Check Now (located in the lower left-hand corner of the window) to check for the latest updates. The link in the Action column indicates whether an update is available. If a version is available, the **Download** link displays.
3. Click **Download** to download a selected version. After successful download, the link in the **Action** column changes from **Download** to **Install**.
4. Click **Install** and select the devices on which you want to install the update. When the installation completes, a check mark displays in the **Currently Installed** column.

STEP 5 | Deploy software updates to selected firewalls.

*If your firewalls are configured in HA, make sure to clear the **Group HA Peers** check box and upgrade one HA peer at a time.*

1. Select **Panorama** > **Device Deployment** > **Software**.
2. Check for the latest updates. Click **Check Now** (located in the lower left-hand corner of the window) to check for the latest updates. The link in the **Action** column indicates whether an update is available.
3. Review the **File Name** and click **Download**. Verify that the software versions that you download match the firewall models deployed on your network. After successful download, the link in the **Action** column changes from **Download** to **Install**.
4. Click **Install** and select the devices on which you want to install the software version.
5. Select **Reboot device after install**, and click **OK**.
6. If you have devices configured in HA, clear the **Group HA Peers** check box and upgrade one HA peer at a time.

STEP 6 | Verify the software and Content Release version running on each managed device.

1. Select **Panorama** > **Managed Devices**.
2. Locate the device(s) and review the content and software versions on the table.

 ©2024 Palo Alto Networks, Inc.

Upgrade the VM-Series for NSX Without Disrupting Traffic

Use the following procedure to upgrade the PAN-OS version of the VM-Series firewalls in your VMware NSX environment. This procedure allows you to perform the PAN-OS upgrade without disrupting traffic by migrating VMs to different ESXi hosts.

STEP 1 | Review the VM-Series for VMware NSX upgrade paths.

STEP 2 | Save a backup of the current configuration file on each managed firewall that you plan to upgrade.

Although the firewall will automatically create a backup of the configuration, it is a best practice to create a backup prior to upgrade and store it externally.

1. Select **Device** > **Setup** > **Operations** and click **Export Panorama and devices config bundle**. This option is used to manually generate and export the latest version of the configuration backup of Panorama and of each managed device.
2. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the upgrade.

STEP 3 | Check the Release Notes to verify the Content Release version required for the PAN-OS version.

The firewalls you plan to upgrade must be running the Content Release version required for the PAN-OS version.

1. Select **Panorama** > **Device Deployment** > **Dynamic Updates**.
2. Check for the latest updates. Click Check Now (located in the lower left-hand corner of the window) to check for the latest updates. The link in the Action column indicates whether an update is available. If a version is available, the **Download** link displays.
3. Click **Download** to download a selected version. After successful download, the link in the **Action** column changes from **Download** to **Install**.
4. Click **Install** and select the devices on which you want to install the update. When the installation completes, a check mark displays in the **Currently Installed** column.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Download the PAN-OS image to all VM-Series firewalls in the cluster.

1. Login to Panorama.
2. Select **Panorama** > **Device Deployment** > **Software**.
3. Click **Refresh** to view the latest software release and also review the **Release Notes** to view a description of the changes in a release and to view the migration path to install the software.
4. Click **Download** to retrieve the software then click **Install**.

   Do not reboot the VM-Series firewalls after installing the new software image.

5. Select the managed devices to be upgraded.
6. Clear the **Reboot device after install** check box.

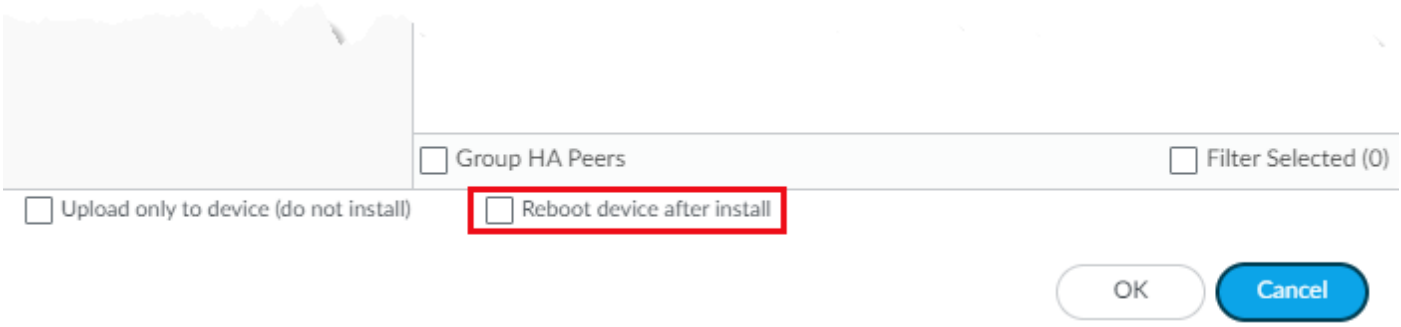

7. Click **OK**.

STEP 5 | Upgrade the VM-Series firewall on the first ESXi host in the cluster.

1. Login to vCenter.
2. Select **Hosts and Clusters**.
3. Right-click the host and select **Maintenance Mode** > **Enter Maintenance Mode**.
4. Migrate (automatically or manually) all VMs, except the VM-Series firewall, off of the host.
5. Power off the VM-Series firewall. This should happen automatically upon entering maintenance mode on the host.
6. (Optional) Assign additional CPUs or memory to the VM-Series firewall before continuing with the upgrade process.

   Verify that enough hardware resources are available to the VM-Series firewall. Refer to the VM-Series System Requirements to see the new resource requirements for each VM-Series model.
7. Right-click the host and select **Maintenance Mode** > **Exit Maintenance Mode**. Exiting maintenance mode causes the NSX ESX Agent Manager (EAM) to power on the VM-Series firewall. The firewall reboots with the new PAN-OS version.
8. Migrate (automatically or manually) all VMs back to the original host.

STEP 6 | Repeat this process for each VM-Series firewall on each ESXi host.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Verify the software and Content Release version running on each managed device.

1. Select **Panorama** > **Managed Devices**.
2. Locate the device(s) and review the content and software versions on the table.

Upgrade the VM-Series Model

The VM-Series firewall licensing process uses the UUID and CPU ID to generate a unique serial number for a firewall. When you generate a license, it is mapped to a specific instance of the VM-Series firewall and cannot be modified. See VM-Series Firewall Licensing for an overview of the license types and consumption models.

Use the instructions in this section if you want to:

- Migrate from a evaluation capacity license to a production license.
- Upgrade the model to increase capacity. For example, you want to upgrade from the VM-100 to the VM-300 model, or change the instance type or VM type for your hypervisor.
  - Upgrading capacity restarts some critical processes on the firewall.
  - If your firewall is licensed with the BYOL option, you must deactivate your VM before you change the model. If you do not, your license is invalidated because a new UUID and CPU ID are assigned when the model changes. After you complete the capacity upgrade, a license for the new capacity must be applied.
  - If you have an HA configuration, see Upgrade the VM-Series Model in an HA Pair.

You can use a VM-Series Model, but fund the license with Software NGFW credits. For that procedure, see

STEP 1 | Allocate additional hardware resources to your VM-Series firewall.

Before initiating the capacity upgrade, you must verify that enough hardware resources are available to the VM-Series firewall to support the new capacity. The process for assigning additional hardware resources differs on each hypervisor.

To check the hardware requirements for your new VM-Series model, see VM-Series Models.

Although the capacity upgrade does not require a reboot of the VM-Series firewall, you need to power down the virtual machine to change the hardware allocation.

STEP 2 | Retrieve the license deactivation API key from the Customer Support portal.

1. Log in to the Customer Support Portal.

   Make sure that you are using the same account that you used to register the initial license.

2. From the menu on the left, select **Assets** > **API Key Management**.

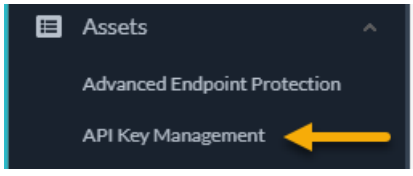

3. Copy the API key.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | On the firewall, use the CLI to install the API key copied in the previous step.

```
request license api-key set key <key>
```

STEP 4 | Enable the firewall to **Verify Update Server identity** on **Device** > **Setup** > **Service**.

STEP 5 | **Commit** your changes. Ensure that you have a locally-configured user on the firewall. Panorama pushed users might not be available after the deactivation if the configuration exceeds the non-licensed PA-VM objects limit.

STEP 6 | Upgrade the capacity.

Select **Device** > **Licenses** > **Upgrade VM Capacity** and then activate your licenses and subscriptions in one of the following ways:

- **Retrieve license keys from license server**—Use this option if you activated your license on the Customer Support portal.
- **Manually upload license key**—Use this option if your firewall does not have connectivity to the Customer Support portal. In this case, you must download a license key file from the support site on an Internet connected computer and then upload to the firewall.
- **Use an authorization code**—Use this option to upgrade the VM-Series capacity using an authorization code for licenses that have not been previously activated on the support portal. When prompted, enter the **Authorization Code** and then click **OK**.

STEP 7 | Verify that your firewall is licensed successfully.

On the **Device** > **Licenses** page, verify that the license was successfully activated.

Upgrade the VM-Series Model in an HA Pair

Upgrading the VM-Series firewall allows you to increase the capacity on the firewall. Capacity is defined in terms of the number of sessions, rules, security zones, address objects, IPSec VPN tunnels, and SSL VPN tunnels that the VM-Series firewall is optimized to handle. When you apply a new capacity license on the VM-Series firewall, the model number and the associated capacities are implemented on the firewall.

Verify the VM-Series system requirements for your firewall model before you upgrade. If your firewall has less than 5.5GB memory, the capacity (number of sessions, rules, security zones, address objects, etc) on the firewall will be limited to that of the VM-50 Lite.

This process is similar to that of upgrading a pair of hardware-based firewalls that are in an HA configuration. During the capacity upgrade process, session synchronization continues, if you have it enabled. To avoid downtime when upgrading firewalls that are in a high availability (HA) configuration, update one HA peer at a time.

Do not make configuration change to the firewalls during the upgrade process. During the upgrade process, configuration sync is automatically disabled when a capacity mismatch is detected and is then re-enabled when both HA peers have matching capacity licenses.

If the firewalls in the HA pair have different major software versions (such as 9.1 and 9.0) and different capacities, both devices will enter the Suspended HA state. Therefore, it is recommended that you make sure both firewalls are running the same version of PAN-OS before upgrading capacity.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Upgrade the capacity license on the passive firewall.

Follow the procedure to Upgrade the VM-Series Model.

The new VM-Series model displays on the dashboard after some processes restart on this passive peer. This upgraded peer is now is a non-functional state because of the capacity mismatch with its active peer.

If you have enabled session synchronization, verify that sessions are synchronized across HA peers before you continue to the next step. To verify session synchronization, run the **`show high-availability interface ha2`** command and make sure that the Hardware Interface counters on the CPU table are increasing as follows:

- In an active/passive configuration, only the active peer show packets transmitted and the passive device will only show packets received.

  If you have enabled HA2 keep-alive, the hardware interface counters on the passive peer will show both transmit and receive packets. This occurs because HA2 keep-alive is bidirectional which means that both peers transmit HA2 keep-alive packets.
- In an active/active configuration, you will see packets received and packets transmitted on both peers.

STEP 2 | Upgrade the capacity license on the active firewall.

Follow the procedure to Upgrade the VM-Series Model.

The new VM-Series model displays on the dashboard after the critical processes restart. The passive firewall becomes active, and this peer (previously active firewall) moves from the initial state to becoming the passive peer in the HA pair.

Downgrade a VM-Series Firewall to a Previous Release

Use the following workflow to restore the configuration that was running before you upgraded to a different feature release. Any changes made since the upgrade are lost. Therefore, it is important to back up your current configuration so you can restore those changes when you return to the newer release.

Use the following procedure to downgrade to a previous release.

STEP 1 | Save a backup of the current configuration file.

Although the firewall automatically creates a backup of the configuration, it is a best practice to create a backup before you upgrade and store it externally.

1. **Export named configuration snapshot** (**Device** > **Setup** > **Operations**).
2. Select the XML file that contains your running configuration (for example, **running-config.xml**) and click **OK** to export the configuration file.
3. Save the exported file to a location external to the firewall. You can use this backup to restore the configuration if you have problems with the downgrade.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Install the previous feature release image.

Autosave versions are created when you upgrade to a new release.

1. **Check Now** (**Device** > **Software**) for available images.
2. Locate the image to which you want to downgrade. If the image is not already downloaded, then **Download** it.
3. After the download completes, **Install** the image.
4. **Select a Config File for Downgrading**, which the firewall will load after you reboot the device. In most cases, you should select the configuration that was saved automatically when you upgraded from the release to which you are now downgrading. For example, if you are running PAN-OS 9.1 and are downgrading to PAN-OS 9.0.3, select `autosave-9.0.3`.
5. After the installation completes successfully, reboot using one of the following methods:
   - If you are prompted to reboot, click **Yes**.
   - If you are not prompted to reboot, go to Device Operations (**Device** > **Setup** > **Operations**) and **Reboot Device**.

 ©2024 Palo Alto Networks, Inc.

VM-Series Plugin

The VM-Series firewalls include the VM-Series plugin, a built-in plugin architecture for integration with public cloud providers or private cloud hypervisors. The VM-Series plugin can be manually upgraded independent of PAN-OS, enabling Palo Alto Networks® to accelerate the release of new features, fixes, or integrations with new cloud providers or hypervisors.

The VM-Series plugin enables you to manage cloud-specific interactions between the VM-Series firewalls and the supported public cloud platforms—AWS, GCP, and Azure. The plugin enables publishing custom metrics to cloud monitoring services (such as AWS CloudWatch), bootstrapping, configuring user credential provisioning information from public cloud environments, and seamless updates for cloud libraries or agents on PAN-OS.

The VM-Series plugin does not manage capabilities that are common to both VM-Series firewalls and hardware-based firewalls. For example, VM Monitoring is not part of the VM-Series plugin because it is a core PAN-OS feature that helps you enforce policy consistently on your virtual machine workloads from both VM-Series firewalls and hardware-based firewalls.

The VM-Series plugin does not manage Panorama plugins*. For the difference between the VM-Series plugin and Panorama plugins, see* VM-Series Plugin and Panorama Plugins*.*

The VM-Series plugin is a built-in component that can be upgraded or downgraded, but not removed. Each PAN-OS release includes a specific VM-Series plugin version that corresponds to the PAN-OS software version. When you downgrade to an earlier PAN-OS software version, the plugin version is downgraded to the version compatible with the PAN-OS version. You can upgrade or downgrade the VM-Series plugin locally on the virtual firewall, or manage the plugin version centrally from Panorama.

To enable Panorama to manage the VM-Series plugin version itself, or cloud-specific metrics publishing your managed firewalls, you must manually install the VM-Series plugin on Panorama as described in Panorama Plugins.

- Configure the VM-Series Plugin on the Firewall
- Upgrade the VM-Series Plugin

Configure the VM-Series Plugin on the Firewall

Select **Device** > **VM-Series** to configure the plugin integration for the cloud provider on which this instance of the VM-Series firewall is deployed.

 ©2024 Palo Alto Networks, Inc.

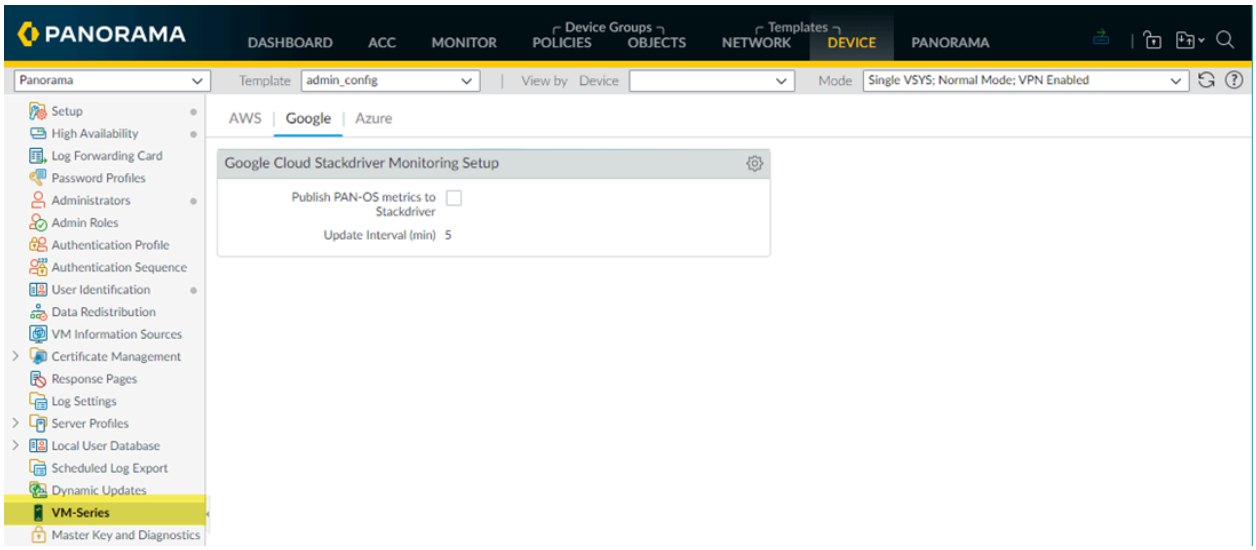

If your firewall is deployed on a hypervisor or cloud without a public interface (for example, VMware ESXi), the tab is named VM-Series and displays a general message.

Upgrade the VM-Series Plugin

When a plugin update is released independent of PAN-OS, you can independently upgrade the plugin version from your VM-Series firewall (like software or content updates) or from a bootstrap file.

Each plugin version provides PAN-OS compatibility information and includes new features or bug fixes for one or more cloud environments.

STEP 1 | Before upgrading, check the latest Release Notes for details on whether a new VM-Series plugin affects your environment.

For example, suppose a new VM-Series plugin version only includes AWS features. To take advantage of the new features, you must update the plugin on your VM-Series firewall instances on AWS.

Do not install an upgrade that does not apply to your environment.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Log in to the VM-Series firewall and check the dashboard to view the plugin version.

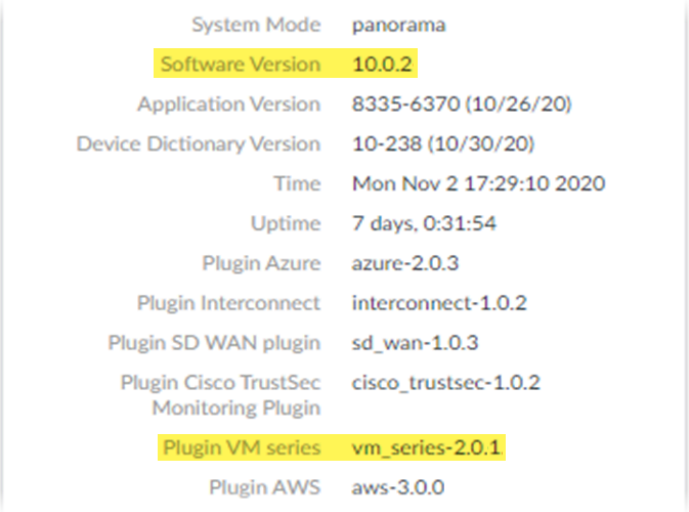

STEP 3 | Select **Panorama** > **Plugins** and type **`vm_series`** in the search field.

Select **Check Now** to view the available versions.

STEP 4 | Choose a VM-Series plugin version and click **Download**.

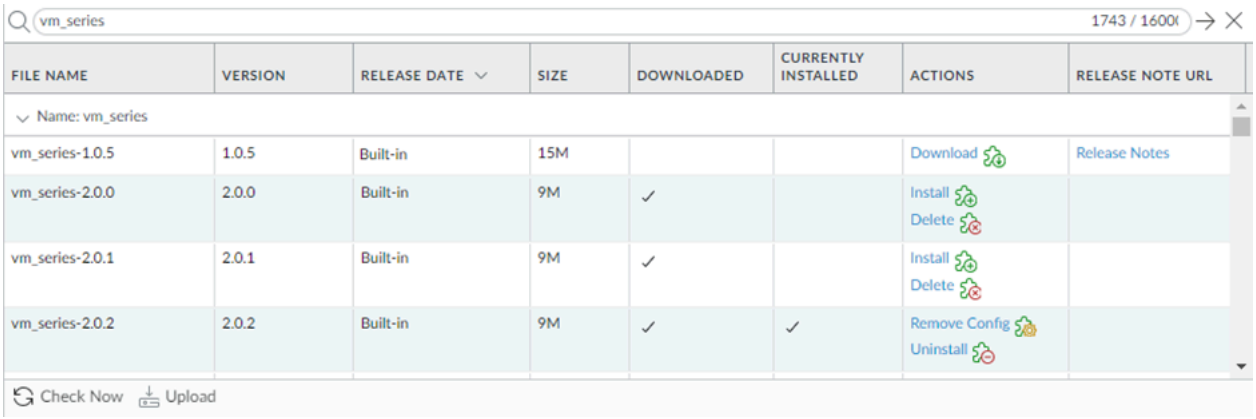

STEP 5 | When the download finishes, click **Install** in the **Actions** column.

The firewall automatically uninstalls the previously installed version of the plugin.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | View the **Dashboard** to verify that the plugin upgraded successfully.

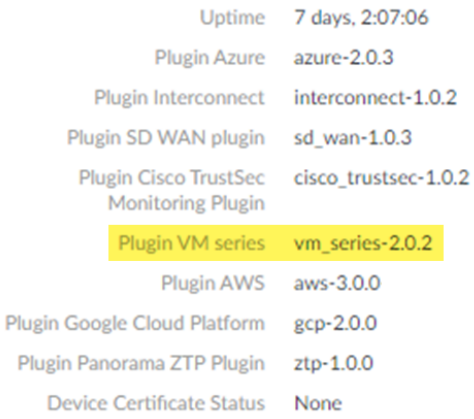

 ©2024 Palo Alto Networks, Inc.

Enable Jumbo Frames on the VM-Series Firewall

By default, the maximum transmission unit (MTU) size for packets sent on a Layer 3 interface is 1500 bytes. This size can be manually set to any size from 512 to 1500 bytes on a per-interface basis. Some configurations require Ethernet frames with an MTU value greater than 1500 bytes. These are called jumbo frames.

To use jumbo frames on a firewall you must specifically enable jumbo frames at the global level. When this is enabled, the default MTU size for all Layer 3 interfaces is set to a value of 9192 bytes. This default value can then be set to any value in the range of 512 to 9216 bytes.

After setting a global jumbo frame size it becomes the default value for all Layer 3 interfaces that have not explicitly had an MTU value set at the interface configuration level. This can become a problem if you only want to exchange jumbo frames on some interfaces. In these situations, you must set the MTU value at every Layer 3 interface that you do not want to use the default value.

The following procedure describes how to enable jumbo frames on a firewall, set the default MTU value for all Layer 3 interfaces and to then set a different value for a specific interface.

VM-Series firewall instances deployed with multiple NUMA nodes, come up in packet MMAP mode when jumbo frame support is enabled. You must disable jumbo frame support to use DPDK on VM-Series firewall instance deployed with multiple NUMA nodes.

STEP 1 | Enable jumbo frames and set a default global MTU value.

1. Select **Device** > **Setup** > **Session** and edit the Session Settings section.
2. Select **Enable Jumbo Frame**.
3. Enter a value for **Global MTU**.

   The default value is 9192. The range of acceptable values is: 512 - 9216.
4. Click **OK**.

   A message is displayed that informs you that enabling or disabling Jumbo Frame mode requires a reboot and that Layer 3 interfaces inherit the **Global MTU** value.
5. Click **Yes**.

   A message is displayed to inform you that Jumbo Frame support has been enabled and reminds you that a device reboot is required for this change to be activated.
6. Click **OK**.
7. Click **Commit**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Set the MTU value for a Layer 3 interface and reboot the firewall.

The value set for the interface overrides the global MTU value.

1. Select **Network** > **Interfaces**.
2. Select an interface of the Layer3 **Interface type**.
3. Select **Advanced** > **Other Info**.
4. Enter a value for **MTU**.

   The default value is 9192. The range of acceptable values is: 512 - 9216.
5. Click **OK**.
6. Click **Commit**.
7. Select **Device** > **Setup** > **Operations** and select **Reboot Device**.

 ©2024 Palo Alto Networks, Inc.

Hypervisor Assigned MAC Addresses

By default, the VM-Series firewall uses the MAC address assigned to the physical interface by the host/hypervisor to deploy a VM-Series firewall with Layer 3 interfaces. The firewall can then use the hypervisor assigned MAC address in its ARP responses. This capability allows non-learning switches, such as the VMware vSwitch to forward traffic to the dataplane interface on the firewall without requiring that promiscuous mode be enabled on the vSwitch. If neither promiscuous mode nor the use of hypervisor assigned MAC address is enabled, the host will drop the frame when it detects a mismatch between the destination MAC address for an interface and the host-assigned MAC address.

There is no option to enable or disable the use of hypervisor assigned MAC addresses on AWS and Azure. It is enabled by default for both platforms and cannot be disabled.

If you are deploying the VM-Series firewall in Layer 2, virtual wire, or tap interface modes, you must enable promiscuous mode on the virtual switch to which the firewall is connected. The use of hypervisor assigned MAC address is only relevant for Layer 3 deployments where the firewall is typically the default gateway for the guest virtual machines.

When hypervisor assigned MAC address functionality is enabled on the VM-Series firewall, make note of the following requirements:

- **IPv6 Address on an Interface**—In an active/passive HA configuration (see VM-Series in High Availability), Layer 3 interfaces using IPv6 addresses must not use the EUI-64 generated address as the interface identifier (Interface ID). Because the EUI-64 uses the 48-bit MAC address of the interface to derive the IPv6 address for the interface, the IP address is not static. This results in a change in the IP address for the HA peer when the hardware hosting the VM-Series firewall changes on failover, and leads to an HA failure.
- **Lease on an IP Address**—When the MAC address changes, DHCP client, DHCP relay and PPPoE interfaces might release the IP address because the original IP address lease could terminate.
- **MAC address and Gratuitous ARP**—VM-Series firewalls with hypervisor assigned MAC addresses in a high-availability configuration behave differently than the hardware appliances with respect to MAC addressing. Hardware firewalls use self-generated floating MAC addresses between devices in an HA pair, and the unique MAC address used on each dataplane interface (say eth 1/1) is replaced with a virtual MAC address that is common to the dataplane interface on both HA peers. When you enable the use of the hypervisor assigned MAC address on the VM-Series firewall in HA, the virtual MAC address is not used. The dataplane interface on each HA peer is unique and as specified by the hypervisor.

  Because each dataplane interface has a unique MAC address, when a failover occurs, the now active VM-Series firewall must send a gratuitous ARP so that neighboring devices can learn the updated MAC/IP address pairing. Hence, to enable a stateful failover, the networking devices must not block or ignore gratuitous ARPs; make sure to disable the anti-ARP poisoning feature on the internetworking devices, if required.

Perform the following steps to configure the VM-Series firewall to use the interface MAC addresses provided by the host/hypervisor.

STEP 1 | Select **Device** > **Management** > **Setup** > **General Settings**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Click the **Edit** icon.

STEP 3 | Check the option to **Use Hypervisor Assigned MAC Address**.

When the MAC address change occurs, the firewall generates a system log to record this transition and the interface generates a gratuitous ARP.

STEP 4 | Click **OK**.

STEP 5 | **Commit** the change on the firewall. You do not need to reboot the firewall.

 ©2024 Palo Alto Networks, Inc.

Custom PAN-OS Metrics Published for Monitoring

The firewall natively publishes the following metrics to monitoring systems in the public cloud such as AWS® CloudWatch, Azure® Application Insights, and Google® Stackdriver. These metrics allow you to assess firewall performance and usage patterns so that you can set alarms and take action to automate events such as launching or terminating instances of the VM-Series firewalls. Because these metrics are published through content updates on the firewall, make sure that you have the minimum content release version that is required to enable this capability on your VM-Series firewall.

Metric	Description
Dataplane CPU Utilization (%)	Monitors dataplane CPU usage and measures the traffic load on the firewall.
Dataplane Packet Buffer Utilization (%)	Monitors dataplane buffer usage and measures buffer utilization. If you have a sudden burst in traffic, monitoring your buffer utilization allows you to ensure that the firewall does not deplete the dataplane buffer, which results in dropped packets.
GlobalProtect™ Gateway Active Tunnels	Monitors the number of active GlobalProtect sessions on a firewall deployed as a GlobalProtect gateway. Use this metric if you use this VM-Series firewall as a VPN gateway to secure remote users. Check the datasheet for the maximum number of active tunnels supported for your firewall model.
GlobalProtect Gateway Tunnel Utilization (%)	Monitors the active GlobalProtect tunnels on a gateway and measures tunnel utilization. Use this metric if you use this VM-Series firewall as a VPN gateway to secure remote users.
panSessionConnectionsPerSecond	Monitors the new connection establish rate per second.
panSessionThroughputKbps	Monitors the throughput in Kbps.
panSessionThroughputPps	Monitors the number of packets per second.
Sessions Active	Monitors the total number of sessions that are active on the firewall. An active session is a session that is in the flow lookup table for which packets will be inspected and forwarded, as required by policy.
Session Utilization (%)	Monitors the TCP, UDP, ICMP and SSL sessions that are currently active and the packet rate, new connection

 ©2024 Palo Alto Networks, Inc.

Metric	Description
	establish rate, and firewall throughput to determine session utilization.
SSLProxyUtilization (%)	Monitors the percentage of SSL forward proxy sessions with clients for SSL/TLS decryption.

To publish these metrics, see:

- Enable CloudWatch Monitoring on the VM-Series Firewall
- Enable Azure Application Insights on the VM-Series Firewall
- Enable Google Stackdriver Monitoring on the VM Series Firewall

 ©2024 Palo Alto Networks, Inc.

Interface Used for Accessing External Services on the VM-Series Firewall

To access the Palo Alto Networks servers for retrieving licenses and software and content updates, and for publishing custom PAN-OS metrics or retrieving IP address and tag mapping for monitoring virtual machines in your deployment, the VM-Series firewall uses the management interface except where noted below. To use a dataplane interface instead of the management interface where supported, you must set up a service route that specifies the dataplane interface that the firewall can use to access the server or service.

Access to Server or Service	Interface Used on the VM-Series Firewall
Licensing	Management interface only
Software Updates	Management interface or Service Route
Bootstrapping from a cloud storage location such as AWS S3 bucket, Azure storage file service, or Google storage bucket	Management interface only, including when interfaces are swapped *If your bootstrap.xml file includes license authcodes, you cannot use a service route. To license the firewall, the management interface must be used.*
Publishing PAN-OS metrics to a cloud monitoring service such as AWS CloudWatch, Azure Application Insights or Google Stackdriver	Management interface only, including when interfaces are swapped
VM Monitoring	Management interface or Service Route

 ©2024 Palo Alto Networks, Inc.

PacketMMAP and DPDK Driver Support

Single-root input/output virtualization (SR-IOV) relies on communication between virtual function (VF) drivers on the VM-Series firewall, and physical function (PF) drivers on the host (the hypervisor). The host uses PF drivers to talk to its physical NICs, and the VM-Series firewall uses VF drivers to talk to the PF drivers.

The following diagram is a simple visualization of that concept.

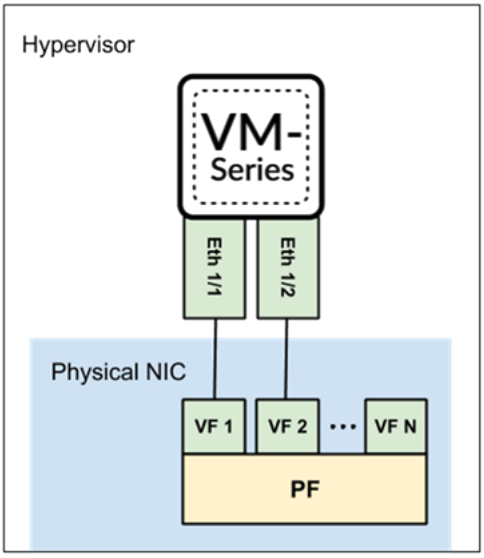

SR-IOV

Why use SR-IOV? SR-IOV is a packet acceleration technology that allows a virtual machine to directly access packets from the NIC. In contrast, when using a virtual switch, the host processes the packets, send the packets through a virtual switch, and then the virtual machine receives its packets.

In the Compatibility Matrix, PacketMMAP Driver Versions lists both the host version and the native driver version on the VM-Series firewall. For example, i40e on the host, and on the firewall, i40e (for PCI-passthrough) and i40evf (for SR-IOV).

For SR-IOV, let's consider a NIC that uses the i40e PF driver. The host communicates with the NIC via the i40e driver. The VM-Series firewall can use its VF driver (i40evf) to directly communicate with the host's PF driver. This allows VM-Series firewall direct access, which improves packet processing speed. To ensure compatibility, install a host PF driver version that is later than the native PF driver version.

PCI-Passthrough

Why does VM-Series firewall have native PF drivers? As mentioned in Options for Attaching VM-Series on the Network, when using PCI-passthrough, the NIC is reserved for the VM-Series firewall, so the host (or other guests on the host) cannot access the NIC. In a PCI-passthrough configuration, the VM-Series firewall uses its native PF driver to communicate directly with the host NIC.

 ©2024 Palo Alto Networks, Inc.

Refer to the PacketMMAP Driver Versions list to determine which PF driver version to install on the host. Install a PF version that is higher than VM-Series firewall native PF driver.

Refer to Enable SR-IOV on ESXi and Enable SR-IOV on KVM for PCI-Passthrough.

DPDK

PAN-OS has two packet processing modes—DPDK (default) and MMAP—and each mode has a corresponding native driver on the VM-Series firewall. For example, if the firewall is in DPDK mode, the firewall uses the DPDK i40evf driver version to communicate with the host's i40e driver (when using SR-IOV). Alternatively, when the firewall is Packet MMAP, it will use a different i40evf driver version to communicate with the host's i40e driver.

You can enable DPDK on the host (the hypervisor), or on the guest (the VM-Series firewall). Enabling both yields the best results.

- Compiling OVS with DPDK is part of enabling DPDK on the host.

  Refer to Configure OVS and DPDK on the Host.
- VM-Series DPDK enables the native DPDK driver on the VM-Series firewall, so DPDK does not need to be enabled on the host, but it is recommended for best performance.

 ©2024 Palo Alto Networks, Inc.

Enable ZRAM on the VM-Series Firewall

If your VM-Series firewall experiences low or out-of memory conditions, you can enable ZRAM to improve memory usage. ZRAM, also called compcache (compressed cache), is a Linux kernel module for creating a compressed block device in RAM. When enabled, ZRAM is used as swap disk and allows for faster I/O of swap because it resides in the RAM.

Complete the following steps to enable ZRAM.

STEP 1 | Log in to the VM-Series CLI.

STEP 2 | Find the total memory on the VM by using the following CLI command.

```
grep pattern “KiB Mem :” mp-log mp-monitor.log
```

```
KiB Mem :  9202656 total,   566504 free,  3475840 used,  5160312
 buff/cache
KiB Mem :  9202656 total,   497112 free,  3481944 used,  5223600
 buff/cache
KiB Mem :  9202656 total,   511744 free,  3466768 used,  5224144
 buff/cache
KiB Mem :  9202656 total,   511668 free,  3466340 used,  5224648
 buff/cache
KiB Mem :  9202656 total,   512124 free,  3465700 used,  5224832
 buff/cache
KiB Mem :  9202656 total,   511436 free,  3465976 used,  5225244
 buff/cache
KiB Mem :  9202656 total,   510984 free,  3465944 used,  5225728
 buff/cache
```

STEP 3 | Convert the above total memory from KB to MB. For example:

9202656 / 1024 = 8987 MB

Take note of the total memory value in MB. You will need this value in the next step.

STEP 4 | Enable ZRAM using the following two CLI commands.

```
debug software kernelcfg zram-swap enable
```

```
debug software kernelcfg zram-swap modify host-mem-threshold <total-
memory-in-MB>
```

STEP 5 | Reboot the VM-Series firewall.

STEP 6 | Verify that ZRAM is enabled.

```
debug software kernelcfg zram-swap show config
```

 ©2024 Palo Alto Networks, Inc.

License the VM-Series Firewall

VM-Series firewall supports two license types (BYOL and PayGo), and two different licensing models—Software Next Generation Firewall Credits (Software NGFW) for flexible configurations that you specify with a deployment profile, and fixed VM-Series Model configurations. Both models also license Security services and other features.

If you are an authorized CSSP partner, see Licenses for Cloud Security Service Providers (CSSPs) for information that pertains to you.

See the following topics for details on creating a support account and managing licenses:

- VM-Series Firewall Licensing
- Create a Support Account
- Serial Number and CPU ID Format for the VM-Series Firewall
- Licensing API
- Install a License API Key
- Use Panorama-Based Software Firewall License Management
- Software NGFW Credits
- VM-Series Models
- What Happens When Licenses Expire?
- Licenses for Cloud Security Service Providers (CSSPs)

VM-Series Firewall Licensing

This chapter compares the following license information:

- License Types: BYOL versus PayGo
- Flexible vCPUs and Fixed Model Licensing: Flexible vCPUs versus fixed models)
- Flexible vCPUs and Fixed Model Deployment: Summary of deployment steps for flexible and fixed models.

License Types

New capacity licenses (non-Software NGFW Credits) are no longer available for purchase. However, you one (1) year renewals for capacity (perpetual and term-based) licenses are available.

Palo Alto Networks currently supports two license types: Bring Your Own License (BYOL) and PAYG (Pay-As-You-Go, also called PayGo).

Type	Description
BYOL	Software NGFW Credits—Available on VM-Series firewalls running all PAN-OS releases. VM-Series firewalls running PAN-OS versions 10.0.4 and later offer advanced features and more flexibility. The flexible license cost is based on the number of vCPUs, the security services you have enabled, and whether you choose to provision Panorama to manage the firewall or act as a log collector. See Software NGFW Credits for a detailed explanation.
BYOL	VM-Series Model licenses—Available for use with all PAN-OS releases. The number of vCPUs is fixed according to your chosen VM-Series model. *Flexible vCPUs, available with PAN-OS 10.0.4 and later, support advanced features and more vCPUs.* The capacity license cost is based on the VM-Series model, the device memory, storage costs, and the support entitlement. Security services and a Panorama deployment to manage your firewalls are additional costs. The capacity license types are: • **VM-Series Enterprise License Agreement (Multi-Model ELA)**—A comprehensive one- or three-year licensing agreement for VM-Series firewalls. An individual license can include a model, security services, a support entitlement, and an optional device management license for Panorama. Multi-Model ELA features a token pool from which you allocate tokens to license VM-Series firewalls. (It is unique to the ELA, and is not the same as the Software NGFW Credits pool.) • Perpetual VM-Series model capacity license with a support entitlement and/or security services bundle 1 or bundle 2.

 ©2024 Palo Alto Networks, Inc.

Type	Description
	• Term firewall capacity license with a support entitlement and your choice of security services.
PayGo	Purchased from a public cloud marketplace (such as AWS, Azure, or GCP), or a Cloud Security Service Provider (CSSP). Available on the PAN-OS version your provider supports. On PAN-OS versions earlier than 9.1.1, PayGo supported only the VM-Series VM-300 model. For PAN-OS 9.1.1 and later PayGo can support fixed Models. The traditional VM models, such as VM-100, VM-300, VM-500, and VM-700 are supported.

Flexible vCPUs and Fixed Model Licensing

What is the difference between flexible vCPU Software NGFW licensing and fixed vCPU VM-Series Model licenses? They charge for different things, and they fund them differently. The following tables provide a quick comparison, and links to greater details.

	Flexible vCPUs	VM-Series Model (Fixed vCPUs)
Description	Cost is based on the number of vCPUs and your chosen Security services. There is no cost for Panorama other than the vCPUs it consumes. You purchase reusable Software NGFW credits that expire at the end of a predetermined term. After activating your credits you can portion them into credit pools. To use your credits, choose a credit profile and create one or more deployment profiles. Choose your own combination of firewall-as-a-platform components: VM-Series vCPUs, security services, virtual Panorama for Management or Dedicated Log Collection, and a support entitlement. All firewalls deployed with a profile are licensed with the same auth code, and you can manage them from the deployment profile.	Cost is based on the VM-Series model capacity license, device memory, and storage. Panorama and Security services are separate purchases. • VM-Series Enterprise License Agreement (Multi-Model ELA)—A comprehensive one- or three-year licensing agreement for VM-Series firewalls. Multi-Model ELA features a token pool from which you allocate tokens to license VM-Series firewalls. • Perpetual VM-Series model capacity license with a support entitlement and/or security services bundle 1 or bundle 2. • Term firewall capacity license with a support entitlement and your choice of security services.
Activation	Requires an activation email. Activation and registration occur automatically.	Requires an activation email and a separate registration step after activation.

 ©2024 Palo Alto Networks, Inc.

	Flexible vCPUs	VM-Series Model (Fixed vCPUs)
Security services	Threat Prevention, DNS Security, GlobalProtect, WildFire, URL Filtering, SD-WAN, DLP, and other services as they become available. When you create your deployment profile you can choose any combination of security services. You can add or remove security services from your profile at any time.	**Bundle 1**: Threat Prevention and premium support entitlement. **Bundle 2**: Threat Prevention, DNS Security, GlobalProtect, WildFire, URL Filtering, SD-WAN, DLP, and premium support entitlement.
PAN-OS version	Up to 32 flexible vCPUs and advanced service options for firewalls running 10.0.4 and later.	You can deploy a VM-Series model (fixed vCPUs) on any PAN-OS version.
Funding	Reusable credits that allow you to consume firewall-as-a-platform components. After you purchase credits you must activate them, associating them to a particular account for your organization. Activated credits fund a credit pool from which you can create a deployment profile. When firewalls are deployed, credits are consumed. When firewalls are deactivated, the credits are released and returned to your credit pool for further use.	• Multi-Model ELA: tokens. • Perpetual VM-Series model capacity license with a support entitlement and/or security services bundle 1 or bundle 2. You determine the configuration at time of purchase. You cannot change the configuration unless you purchase a new license. • Term firewall capacity license with a support entitlement and your choice of security services.
Deployment Configuration	Flexible. A deployment profile can be changed at any time. Changes to the profile propagate to all firewalls that share the deployment profile auth code.	VM-Series model capacity does not change, but if you have an ELA, you an can add Security services. Perpetual and Term licenses are configured and paid for in advance and do not change.
Deployment	After credit activation, create a deployment profile for a specific environment or use case (such as “Protect my NSX Environment”) and configure firewall vCPUs, security services, and an optional virtual Panorama. You can create any number of deployment profiles and customize them at any point in time.	Accept the VM-Series ELA. Deploy and configure the VM-Series firewall. Activate the model license and register the firewall.

 ©2024 Palo Alto Networks, Inc.

	Flexible vCPUs	VM-Series Model (Fixed vCPUs)
	You must have the Customer Support Portal role Credit Administrator (applies to account management only) to activate and manage Software NGFW credits.	
Panorama	When you create a deployment profile you can choose to add Panorama for management, or as a dedicated log collector for firewalls that use a deployment profile. This Panorama can manage firewalls deployed with the deployment profile's shared auth code.	Panorama is a separate expense. A physical or virtual Panorama can be used to for firewall management or for log collection.
Upgrade or Downgrade	If the VM-Series firewall or Panorama has an internet connection, changes to your deployment profile are automatically applied to the firewall. If the firewall does not have an internet connection, manually stop the firewall. In Assets > Software NGFW Credits change the deployment profile, then in the CSP, download the license keys, and transfer them to the VM, obtain the profile from the CSP, transfer it to the VM, restart the VM and apply the license. You do not have to reboot the firewall in either case.	Change to a different model requires a license change and a reboot.

Flexible vCPUs and Fixed Model Deployment

The following checklists compare the deployment processes for Software NGFW credits and the VM-Series Model licensing methods.

Flexible vCPUs	Fixed vCPUs (VM-Series Model)
1. Create a Support Account. 2. Activate Credits. Your organization can have many accounts to represent different cost centers. During registration you associate your credit purchase with an account. 3. Create a Deployment Profile.	1. Create a Support Account. 2. Activate VM-Series Model Licenses. 3. Register the VM-Series Firewall. 4. Deploy the VM-Series firewall on Alibaba, AWS, Azure, Cisco ACI, Cisco CSP, Cisco ENCS, ESXi, Google Cloud Platform, Hyper-V, KVM, OpenStack. Oracle Cloud

 ©2024 Palo Alto Networks, Inc.

Flexible vCPUs	Fixed vCPUs (VM-Series Model)
4. Deploy the VM-Series firewall on Alibaba, AWS, Azure, Cisco ACI, Cisco CSP, Cisco ENCS, ESXi, Google Cloud Platform, Hyper-V, KVM, OpenStack. Oracle Cloud Infrastructure, vCloud Air, NSX-T, or NSX-V 5. Install a Device Certificate on the VM-Series Firewall (for site licenses such as Cortex Data Lake and Auto Focus).	Infrastructure, vCloud Air, NSX-T, or NSX-V. 5. Install a Device Certificate on the VM-Series Firewall (for site licenses such as Cortex Data Lake and Auto Focus).

 ©2024 Palo Alto Networks, Inc.

Create a Support Account

You need a support account to log in to the Customer Support Portal (CSP). You must log in to activate and manage Software NGFW credits, access software updates, or open a case with Palo Alto Networks technical support. Your support account allows you to view and manage all assets—appliances, licenses, and subscriptions—that you have registered with Palo Alto Networks.

For all licensing options, except for usage-based licenses that are currently only available in AWS, you require a support account so that you can download the software package required to install the VM-Series firewall.

If you have an existing support account, you can download and install the VM-Series firewall software, then continue to Register the VM-Series Firewall.

STEP 1 | Go to https://support.paloaltonetworks.com/UserAccount/PreRegister.

STEP 2 | Enter the corporate email address to associate with the support account.

STEP 3 | Choose one of the following options and fill in the details in the user registration form:

For a usage-based license in AWS

1. Click **Register your Amazon Web Services VM-Series Instance**.
2. On the AWS Management Console, find the AWS Instance ID, AWS Product Code, and the AWS Zone in which you deployed the firewall.
3. Fill in the other details.

For all other licenses

1. Click **Register device using Serial Number or Authorization Code**.
2. Enter the capacity auth code and the sales order number or customer ID.
3. Fill in the other details.

STEP 4 | **Submit** the form. You will receive an email with a link to activate your user account.

Complete the steps to activate the account. After your account is verified and the registration is complete, you can log in to the support portal.

 ©2024 Palo Alto Networks, Inc.

Serial Number and CPU ID Format for the VM-Series Firewall

When you launch an instance of the VM-Series firewall, each instance of the firewall is uniquely identified using the CPU ID and serial number of the firewall. The CPU ID format and the serial number include information on the hypervisor and the license type for each instance of the VM-Series firewall.

- With the usage-based licensing model of the VM-Series firewalls, at launch the firewall generates a serial number and CPU ID, and you use these details to Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code).
- With the BYOL model, you register a VM-Series Firewall fixed model or a VM-Series firewall with a flexible license on the Customer Support portal (CSP).
  - For a firewall with direct internet access, you can apply the auth code on the firewall to generate a license file that includes the serial number.
  - For a firewall that is offline, you must use the CSP to input the CPU ID, UUID, and the auth code to generate a license file that includes the serial number. You can then install the license on the firewall.

License Type	Serial Number	CPU ID
BYOL	15 digits, all numeric Example: 0071 **51** 345678909	<Hypervisor>:<ActualCPUID> Example: **ESX**:12345678
PAYG	15 digits, alphanumeric Example: **4** DE0YTAYOGMYYTN	<Hypervisor>:<Instance-ID>:<CloudProductCode>:<CloudRegion> Example: **AWSMP**:1234567890abcdef0:6kxdw3bbmdeda3o6i1ggqt4km:us-west1

 ©2024 Palo Alto Networks, Inc.

Licensing API

Use the licensing API to register auth codes, retrieve licenses attached to an auth code, renew licenses, or deactivate all licenses on a VM-Series firewall. In addition, the licensing API enables you to license firewalls that do not have direct internet access and cannot reach the Palo Alto Networks license server. You can manage licenses manually or automate licensing with a custom script or an orchestration service.

You can also use the Panorama Software Firewall License plugin for licensing tasks, including offline licensing. The plugin requires Panorama 10.0.0 or later with VM-Series plugin 2.0.4 or later, and your managed VM-Series firewalls must be running PAN-OS 9.1.0 or later and VM-Series plugin 2.0.4 or later; the VM-Series firewall for Azure requires VM-Series plugin 2.0.8.

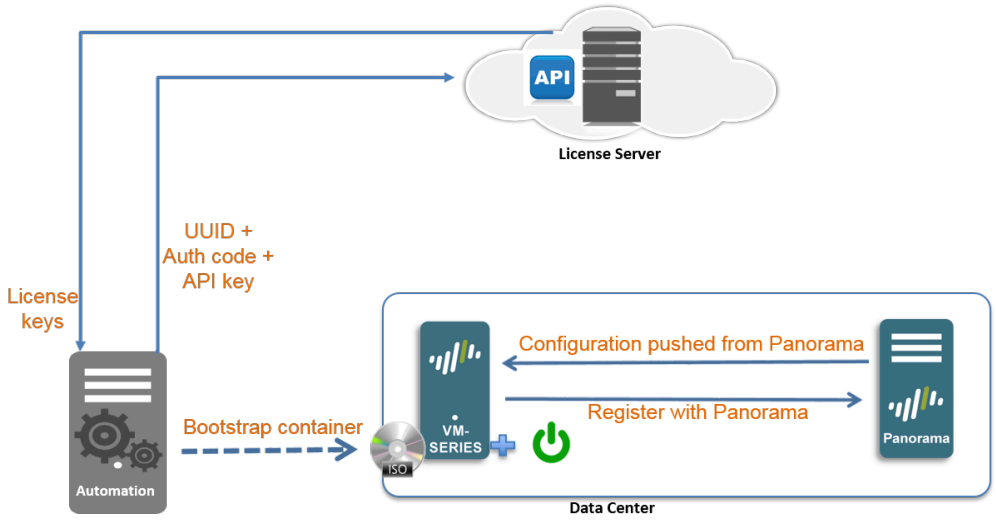

For model-based licenses, he API allows you to view the details of an auth code so that you can track the number of unused licenses attached to an auth-code or auth-code bundle that enables you to license more than one instance of the firewall. An auth-code bundle includes the VM-Series model, subscriptions and support in a single, easy to order format; you can use this bundle multiple times to license VM-Series firewalls as you deploy them.

To use the API, each support account is assigned a unique key. Each API call is a POST request, and the request must include the API key to authenticate the request to the licensing server. When authenticated, the licensing server sends the response in json format (content-type application/json).

- Manage the Licensing API Key
- Use the Licensing API
- Licensing API Error Codes
- Install a License API Key

 ©2024 Palo Alto Networks, Inc.

Manage the Licensing API Key

To get the API key required to use the licensing API, your account must have super user privileges on the support portal. The same key is used to activate and deactivate the license.

The expiration date of the API key is the same date as that of the latest subscription in your support account. If you renew your current subscriptions and need to reset the expiration date of the API key, you can either regenerate a key (and replace the existing key with this new key wherever you've used it) or contact Palo Alto Networks support for help with extending the term of your existing API key.

STEP 1 | Get your Licensing API key.

1. Log in to the Palo Alto Networks Support portal with an account that has super user privileges.
2. Select **Assets** > **API Key Management**.
3. Click **Enable** to view your key and copy it for use. Once you generate a key, the key is enabled until you regenerate or disable it.

STEP 2 | Regenerate or revoke the API key.

1. You can generate a new API key or revoke the use of the key.
   - Click **Regenerate** to generate a new key. If you suspect that an API key may be compromised, you can generate a new key. Regenerating automatically invalidates the old key.
   - Select **Disable** if you no longer plan to use the key. Disabling the API key revokes it.

Use the Licensing API

The base URI for accessing the licensing API is https://api.paloaltonetworks.com/api/license; based on the task you want to perform, for example activate licenses, deactivate licenses, or track license use—the URL will change.

An API request must use the HTTP POST method, and you must include the API key in the apikey HTTP request header and pass the request parameters as URL-encoded form data with `content-type application/x-www-form-urlencoded`.

The API Version is optional and can include the following values—0 or 1. If specified, it must be included in the `version` HTTP request header. The current API version is 1; if you do not specify a version, or specify version 0, the request uses the current API version.

All API responses are represented in json.

Before you begin, get your Licensing API key and copy it to your local drive. This is required before you can perform any of the following tasks:

- Activate Licenses
- Deactivate Licenses
- Track License Usage

Activate Licenses

Header: apikey

 ©2024 Palo Alto Networks, Inc.

Parameters: `uuid`, `cpuid`, `authCode`, `memory`, `serialNumber`, and `vCPU`

URL: **`https://api.paloaltonetworks.com/api/license/activate`**

The parameters uuid, cpuid, authCode, and serialNumber apply for all VM-Series licenses, regardless of PAN-OS version.

The optional parameters `memory` and `vCPU` only apply for Flexible vCPUs (PAN-OS 10.0.4 and later).

- For the initial license activation, provide the parameters in the API request. For example:

```
curl -i -H
 "apikey:a103e3065360acc5e01666fb9335964fcfe668100666db6f3ff43d4544de0###"
--data-urlencode cpuid=AWS:57060500FFFBE###
--data-urlencode uuid=EC2278FF-F0CB-45E2-343B-E97984BAC###
--data-urlencode authCode=D3521###
--data-urlencode vcpu=4
--data-urlencode memory=8388608
https://api.paloaltonetworks.com/api/license/activate
```

- If you did not save the license keys or had network connection issues during initial license activation, you can retrieve the license(s) for a firewall that you have previously activated.

  In the API request, provide the `cpuid` and `uuid`, or, provide the `serialNumber` of the firewall.

Sample request for initial license activation using Curl:

```
curl -i -H "apikey:$APIKEY" --data-urlencode cpuid=51060400FFFBAB1F
 --data-urlencode uuid=564D0E5F-3F22-5FAD-DA58-47352C6229FF --data-
urlencode authCode=I7115398 https://api.paloaltonetworks.com/api/
license/activate
```

Sample API response:

```
[{"lfidField":"13365773","partidField":"PAN-SVC-PREM-
VM-300","featureFi--eld":
"Premium","feature_descField":"24 x 7 phone support; advanced
 replacement hardware service","keyField":"m4iZEL1t3n6Oa
+6ll1L7itDZTphYw48N1AMOZXutDgExC5f5pOA52+Qg1jmAxanB
\nKOyat4FJI4k2hWiBYz9cONuKoiaNOtAGhJvAuZmYgqAZejKueWrTzCuLrwxI/iEw
\nkRGR3cYG+j6o84RitR937m2iOk2v9o8RSfLVilgX28nqmcO8LcAnTqbrRWdFtwVk
\nluz47AUMXauuqwpMipouQYjk0ZL7fTHHslhyL7yFjCyxBoYXOt3JiqQ0OCDdBdDI
\n91RkVPylEwTKgSXm3xpzbmC2ciUR5b235gyqdyW8eQXKvaThuR8YyHr1Pdw/lAjs
\npyyIVFa6FufPacfB2RHApQ==\n","auth_codeField":"","errmsgField":null,
"typeField":"SUP","regDateField":"2016-06-03T08:18:41","startDateField":"5/29/2
"vm_capacityField":null,"uuidField":null,"cpuidField":null,"mac_baseField":null
"mac_countField":null,"drrField":null,"expirationField":"8/29/2016
 12:00:00 AM","PropertyChanged":null},
{"lfidField":"13365774","partidField":"PAN-VM-300-TP",
"featureField":"Threat Prevention","feature_descField":"Threat
 Prevention","keyField":"NqaXoaFG+9qj0t9Vu7FBMizDArj
+pmFaQEd6I2OqfBfAibXrvuoFKeXX/K2yXtrl\n2qJhNq3kwXBDxn181z3nrUOsQd/
eW68dyp4jb1MfAwEM8mlnCyLhDRM3EE+umS4b\ndZBRH5AQjPoaON7xZ46VMFovOR
+asOUJXTptS/Eu1bLAI7PBp3+nm04dYTF9O50O
\ndey1jmGoiBZ9wBkesvukg3dVZ7gxppDvz14+wekYEJqPfM0NZyxsC5dnoxg9pciF
```

 ©2024 Palo Alto Networks, Inc.

```
\ncFelhnTYlma1lXrCqjJcFdniHRwO0RE9CIKWe0g2HGo1uo2eq1XMxL9mE5t025im
\nblMnhL06smrCdtXmb4jjtg==\n","auth_codeField":"",
"errmsgField":null,"typeField":"SUB","regDateField":"2016-06-03T08:18:41",
"startDateField":"5/29/2016","vm_capacityField":null,"uuidField":null,
"cpuidField":null,"mac_baseField":null,"mac_countField":null,"drrField":null,
"expirationField":"8/29/2016 12:00:00 AM","PropertyChanged":null}
...<truncated>
```


The feature_Field in the response indicates the type of key that follows in the keyField. Copy each key to a text file and save it with the .key extension. Because the key is in json format, it does not have newlines. Make sure to convert it to newlines if your parser requires them. Make sure to name each key appropriately and save it to the /license folder of the bootstrap package. For example, include the authcode with the type of key to name it as I3306691_1pa-vm.key (for the capacity license key), I3306691_1threat.key (for the Threat Prevention license key), I3306691_1wildfire.key (for the WildFire subscription license key).

Sample API request for retrieving previously activated licenses using Curl:

```
curl -i -H "apikey:$APIKEY" --data-urlencode
 serialNumber=007200006142 https://api.paloaltonetworks.com/api/
license/activate
```

Sample API response:


```
[{"lfidField":"13365773","partidField":"PAN-SVC-PREM-
VM-300","featureField":
"Premium","feature_descField":"24 x 7 phone support; advanced
 replacement hardware service","keyField":"m4iZEL1t3n6Oa
+6ll1L7itDZTphYw48N1AMOZXutDgExC5f5pOA52+Qg1jmAxanB
\nKOyat4FJI4k2hWiBYz9cONuKoiaNOtAGhJvAuZmYgqAZejKueWrTzCuLrwxI/iEw
\nkRGR3cYG+j6o84RitR937m2iOk2v9o8RSfLVilgX28nqmcO8LcAnTqbrRWdFtwVk
\nluz47AUMXauuqwpMipouQYjk0ZL7fTHHslhyL7yFjCyxBoYXOt3JiqQ0OCDdBdDI
\n91RkVPylEwTKgSXm3xpzbmC2ciUR5b235gyqdyW8eQXKvaThuR8YyHr1Pdw/lAjs
\npyyIVFa6FufPacfB2RHApQ==\n","auth_codeField":"","errmsgField":null,
"typeField":"SUP","regDateField":"2016-06-03T08:18:41","startDateField":"5/29/2
"vm_capacityField":null,"uuidField":null,"cpuidField":null,"mac_baseField":null
"mac_countField":null,"drrField":null, "expirationField":"8/29/2016
 12:00:00 AM","PropertyChanged":null},
{"lfidField":"13365774","partidField":"PAN-VM-300-TP",
"featureField":"Threat Prevention","feature_descField":
"Threat Prevention","keyField":
"NqaXoaFG+9qj0t9Vu7FBMizDArj+pmFaQEd6I2OqfBfAibXrvuoFKeXX/K2yXtrl
\n2qJhNq3kwXBDxn181z3nrUOsQd/eW68dyp4jb1MfAwEM8mlnCyLhDRM3EE+umS4b
\ndZBRH5AQjPoaON7xZ46VMFovOR+asOUJXTptS/Eu1bLAI7PBp3+nm04dYTF9O50O
\ndey1jmGoiBZ9wBkesvukg3dVZ7gxppDvz14+wekYEJqPfM0NZyxsC5dnoxg9pciF
\ncFelhnTYlma1lXrCqjJcFdniHRwO0RE9CIKWe0g2HGo1uo2eq1XMxL9mE5t025im
\nblMnhL06smrCdtXmb4jjtg==
\n","auth_codeField":"","errmsgField":null,"typeField":"SUB",
"regDateField":"2016-06-03T08:18:41","startDateField":"5/29/2016","vm_capacityF
:null,"uuidField":null,"cpuidField":null,"mac_baseField":null,
"mac_countField":null,"drrField":null,"expirationField":"8/29/2016
 12:00:00 AM","PropertyChanged":null}
...<truncated>
```


 ©2024 Palo Alto Networks, Inc.

Deactivate Licenses

URL: **`https://api.paloaltonetworks.com/api/license/deactivate`**

Parameters: encryptedToken

To deactivate the license(s) on a firewall that does not have direct internet access, you must generate the license token file locally on the firewall and then use this token file in the API request. For details on generating the license token file, see Deactivate VM, or Deactivate License (Software NGFW Credits) and Deactivate a Feature License or Subscription Using the CLI.

Header: apikey

Request: https://api.paloaltonetworks.com/api/license/deactivate?encryptedtoken@<token>

Sample API request for license deactivation using Curl:

```
curl -i -H "apikey:$APIKEY" --data-urlencode
 encryptedtoken@dact_lic.05022016.100036.tok https://
api.paloaltonetworks.com/api/license/deactivate
```

Sample API response:

```
[{"serialNumField":"007200006150","featureNameField":"","issueDateField":"",
"successField":"Y","errorField":null,"isBundleField":null,"PropertyChanged":nul
{"serialNumField":"007200006150","featureNameField":"","issueDateField":"",
"successField":"Y","errorField":null,"isBundleField":null,"PropertyChanged":nul
{"serialNumField":"007200006150","featureNameField":"","issueDateField":"",
"successField":"Y","errorField":null,"isBundleField":null,"PropertyChanged":nul
{"serialNumField":"007200006150","featureNameField":"","issueDateField":"",
"successField":"Y","errorField":null,"isBundleField":null,"PropertyChanged":nul
{"serialNumField":"007200006150","featureNameField":"","issueDateField":"",
"successField":"Y","errorField":null,"isBundleField":null,"PropertyChanged":nul
{"serialNumField":"007200006150","featureNameField":"","issueDateField":"",
"successField":"Y","errorField":null,"isBundleField":null,
"PropertyChanged":null}]$
```

Track License Usage

URL: **`https://api.paloaltonetworks.com/api/license/get`**

Parameters: authCode

Header: apikey

Request: https://api.paloaltonetworks.com/api/license/get?authCode=<authcode>

Sample API request for tracking license usage using Curl:

```
curl -i -H "apikey:$APIKEY" --data-urlencode authcode=I9875031
 https://api.paloaltonetworks.com/api/license/get
```

Sample API response:

```
HTTP/1.1 200 OK
Date: Thu, 05 May 2016 20:07:16 GMT
Content-Length: 182
```

 ©2024 Palo Alto Networks, Inc.

```
{"AuthCode":"I9875031","UsedCount":4,"TotalVMCount":10,"UsedDeviceDetails":
[{"UUID":"420006BD-113D-081B-F500-2E7811BE80C
9","CPUID":"D7060200FFFBAB1F","SerialNumber":"007200006142"}]}.....
```

Licensing API Error Codes

The HTTP Error Codes that the licensing server returns are as follows:

- 200 Success
- 400 Error
- 401 Invalid API Key
- 500 Server Error

 ©2024 Palo Alto Networks, Inc.

Install a License API Key

The license API key can be used to activate, change, or deactivate a license. This procedure applies for a VM-Series firewall or Panorama deployment.

You must have Super User privileges to retrieve the license API key from the Customer Support Portal and use the CLI to install the key on the firewall or Panorama.

On Panorama, when you install a license API key Panorama pushes the API key to its managed devices. If the managed device has an API key installed, Panorama overwrites the old API key with the new one.

STEP 1 | Retrieve the license API key from the Customer Support Portal.

1. Log in to the Customer Support Portal.
2. Select **Assets** > **Licensing API**.
3. Copy the API key.

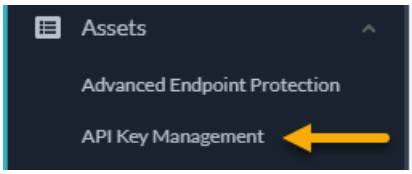

STEP 2 | Use the CLI to install the API key copied in the previous step. Paste the key into the request:

```
request license api-key set key <key>
```

STEP 3 | (optional) To replace a license deactivation API key, use the following CLI command to delete an installed API key.

```
request license api-key delete
```

If you delete the API key you must install another license deactivation key before you can deactivate licenses.

 ©2024 Palo Alto Networks, Inc.

Use Panorama-Based Software Firewall License Management

The Panorama Software Firewall License plugin allows you to automatically license a VM-Series firewall when it connects to Panorama. If your VM-Series firewalls are located in the perimeter of your deployment and do not have connectivity to the Palo Alto Networks licensing server, the Software Firewall License plugin simplifies the license activation process by using Panorama to license the VM-Series firewall.

Additionally, the Software Firewall License plugin simplifies the license activation and deactivation of VM-Series firewalls in environments that use auto-scaling and automation to deploy and delete firewalls to address changes in the cloud.

Pay-as-you-go (PAYG) licenses are not supported for use with this plugin.

Do not use the Software Firewall License plugin to license the VM-Series firewall for VMware NSX. The Panorama plugin for VMware NSX automatically licenses VM-Series firewalls deployed in NSX and NSX-T

.

Also, do not use this plugin to license firewalls deployed in device groups that include instances of the VM-Series firewall deployed in NSX-T.

To install the Panorama Software Firewall License plugin, you must be using Panorama 10.0.0 or later and VM-Series plugin 2.0.4 or later. Your VM-Series firewalls must be running PAN-OS 9.1.0 or later.

The VM-Series firewall for Azure requires VM-Series plugin 2.0.7 or later.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away

 ©2024 Palo Alto Networks, Inc.

- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Install the Software Firewall License Plugin for Panorama.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Plugins**.
3. Click **Check Now** to get the list of available plugins.
4. Search for sw_fw_license to locate the plugin.
5. Select **Download** and **Install** the Software Licensing plugin.

After you successfully install, Panorama refreshes and the Software Licensing plugin displays on the **Panorama** tab.

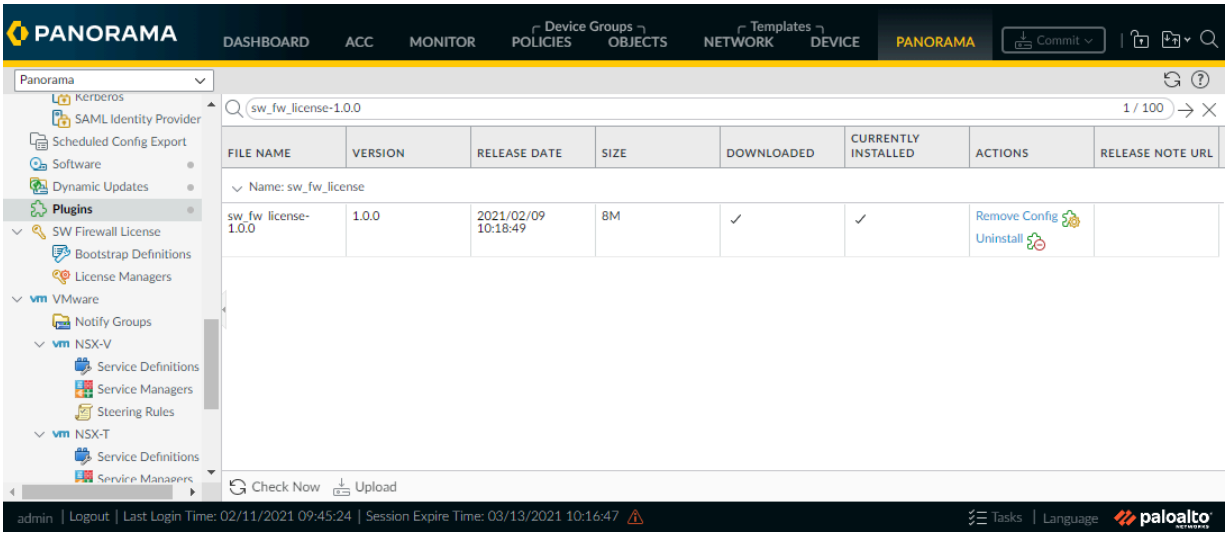

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure a Bootstrap Definition.

1. Select **Panorama** > **SW Firewall License** > **Bootstrap Definitions**.
2. Click **Add**.
3. Enter a descriptive **Name** to identify the Bootstrap Definition.
4. (Optional) Enter a **Description** of the Bootstrap Definition.
5. Enter the **Auth Code** that Panorama will use to license the VM-Series firewall when it connects to Panorama.
6. Click **OK**.

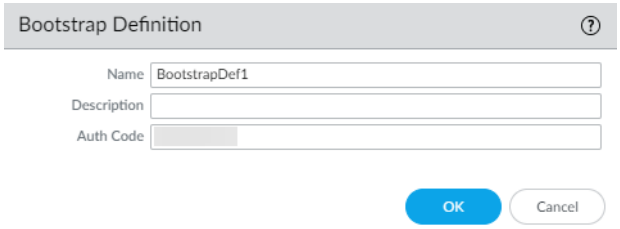

STEP 3 | Configure a License Manager.

1. Select **Panorama** > **SW Firewall License** > **License Managers**.
2. Click **Add**.
3. Enter a descriptive **Name** to identify the License Manager.
4. (Optional) Enter a **Description** of the License Manager.
5. Select a **Device Group** from the drop-down. When a VM-Series firewall bootstrapped using the license manager connects to Panorama, it is placed in the specified device group.
6. Select a **Template Stack** from the drop-down. When a VM-Series firewall bootstrapped using the license manager connects to Panorama, it is placed in the specified template Stack.
7. In the **Auto Deactivate** field, specify the amount of time, in hours, that Panorama waits before deactivating the license of a disconnected VM-Series firewall. When you select **Never**, Panorama does not deactivate a disconnected VM-Series firewall. Auto

 ©2024 Palo Alto Networks, Inc.

Deactivate is set to Never by default. You can set the deactivation time, in hours, from one to 24.

When a Auto Deactivate interval, the plugin might deactivate the license of stopped VM-Series firewalls that are managed by the same license manager of firewalls that deactivate as expected.

8. Select a **Bootstrap Definition** from the drop-down. The selected bootstrap definition specifies the auth code used by Panorama to license the VM-Series firewalls associated with the license manager.
9. Click **OK**.
10. **Commit** your changes.

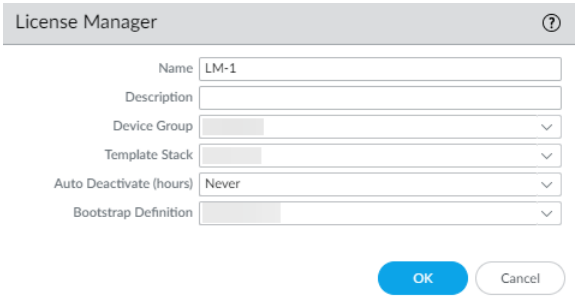

STEP 4 | (Optional) Create an `init-cfg.txt` file for bootstrap the VM-Series firewall. After configuring a license manager, you can copy and paste bootstrap parameters generated by Panorama when deploy your VM-Series firewalls. Depending on your deployment, the parameters displayed might be a subset of those shown in the image below. For example, if your Panorama appliance is deployed in a public cloud, the bootstrap parameters will not include the public IP address for Panorama. In that case, you must manually enter the public IP address in the `init-cfg.txt` file. Panorama will always generate the auth-key

 ©2024 Palo Alto Networks, Inc.

and **`plugin-op-commands=panorama-licensing-mode-on`** for use in your `init-cfg.txt`.

The auth key displayed here is generated by Panorama and used to authenticate the VM-Series firewall connection to Panorama. Additionally, this auth key is used instead of the VM auth key that you might generate on Panorama and add to your init-cfg.txt file.

If you use the auth key displayed here in your init-cfg.txt file, do not use a manually generated VM auth key.

1. Select **Panorama** > **SW Firewall License** > **License Managers**.
2. In the Action column of a given license manager, click **Show Bootstrap Parameters**.
3. Copy the displayed information and paste it into a text editor to create an `init-cfg.txt` file for bootstrapping.
4. Click **Close** when finished.

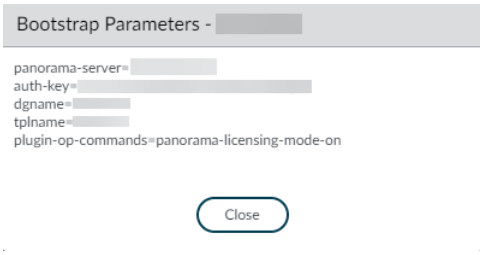

STEP 5 | (Optional) View and deactivate a managed VM-Series firewall. From the **Show Devices** dialogue, you can view the devices associated with a given license manager. You can view the name, serial number, management IP address, connection status, and amount of time Panorama waits to deactivate a disconnected firewall. Additionally, you can manually deactivate the license of managed VM-Series firewall.

1. Select **Panorama** > **SW Firewall License** > **License Managers**.
2. In the Action column of a given license manager, click **Show Devices**.
3. To manually deactivate a connected or disconnected (but not yet deactivated) managed VM-Series firewall, select a one or more listed VM-Series firewalls and click **Deactivate**.

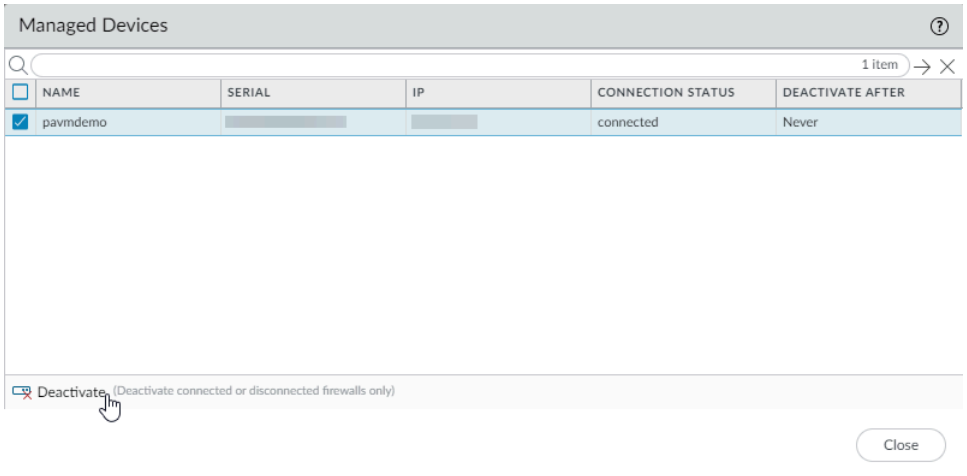

 ©2024 Palo Alto Networks, Inc.

STEP 6 | (Optional) Verify that Panorama has completed the necessary API calls to license connected firewalls.

1. Log in to the Panorama command line interface.
2. Execute the following command.

   ```
   show plugins sw_fw_license panorama-api-requests
   ```

 ©2024 Palo Alto Networks, Inc.

Software NGFW Credits

Software NGFW credits can be used to fund Software NGFWs (VM-Series and CN-Series), Cloud-Delivered Security Services (CDSS), or virtual Panorama appliances in networks with or without internet access (air-gapped networks, for example).

Use Software NGFW Credit Estimator to calculate the number of credits required for your deployment. The Credit Estimator allows you estimate the credits required to license the firewalls and enable the security services needed to secure your deployment.

You create a deployment profile to configure one or more firewalls based on PAN-OS version, the number of vCPUs per firewall, the total number of firewalls supported by the deployment profile, Panorama management or log collection, and security services. All the VMs created with a deployment profile share the same authcode.

- Fixed vCPUs—Compatible with all PAN-OS versions. Based on VM-Series Models and security service bundles. Changing the model or service options requires a new license.
- Flexible vCPUs—Available on PAN-OS 10.0.4 and later. Select a flexible number of vCPUs, and a flexible selection of security services. You can modify the deployment profile to add or decrease the number of vCPUs, add new services as they become available, or remove services. The maximum number of vCPUs for a deployment profile is 32.

Software NGFW credits are term-based. Terms can be defined for any amount of time between 1 and 5 years. Both allocated and unallocated credits expire at the end of the agreed upon term. You can purchase additional credits for a credit pool but the expiration date must be the same as the target pool.

If you have an internet connection to the license server and you stop using a firewall, a security service, or Panorama deployment, the credits allocated to that resource are refunded to the credit pool and can be reallocated to a new resource.

If you do not have an internet connection and cannot connect to the Palo Alto Networks update server (for example, you are in an air-gapped network) you can manage the VM-Series firewall locally from its user interface, or from Panorama. Your administrator must then log in to the Customer Support Portal to return the license token so the funds can be reused.

Use the **Supported Hypervisor** table below and the **Total vCPUs on Dataplane** tables that follow to ensure that you allocate the necessary hardware resources for your chosen number of vCPUs.

Memory Profile	Supported Hypervisors	Minimum Hard Drive
4.5GB, 5.5GB	ESXi, Hyper-V, KVM	32GB (60GB at boot)
6.5GB	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, Cisco ENCS, NSX-T	60GB

 ©2024 Palo Alto Networks, Inc.

Memory Profile	Supported Hypervisors	Minimum Hard Drive
9GB	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, Cisco ENCS, NSX-T	60GB
16GB	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, NSX-T	60GB
56GB	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, NSX-T	60GB

For all memory profiles listed above, the minimum vCPUs is 2. You can enable Lite mode if you choose 4.5GB. Lite mode is an alternative operating mode for environments where resources are limited. See VM-50 Lite Mode for more information.

Lite mode requires minimum 32GB of hard drive space. However, because the VM-Series base image is common for all vCPU combinations, you must allocate 60GB of hard drive space until you license a VM-Series firewall with 4.5GB memory.

To achieve the best performance, all of the required cores should be available on a single CPU socket.

The distribution of vCPUs assigned to the management plane and the dataplane differs depending on the total number of vCPUs assigned to the VM-Series firewall. Refer to the following tables (1-16 vCPUs and 17-32 vCPUs) to find out the number of dataplane vCPUs automatically allocated for a particular combination of memory profile and total vCPUs. For example, with a memory profile of 16GB and a total of 15 vCPUs, 11 are assigned to the dataplane and the remaining 4 are assigned to the management plane.

Memory Profile	Total vCPUs on Dataplane (1-16)														
	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
4.5GB	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
5.5GB	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
6.5GB	1	2	2	2	2	2	2	2	2	2	2	2	2	2	2
9GB	1	2	3	3	4	4	4	4	4	4	4	4	4	4	4
16GB	1	2	3	3	4	5	6	6	7	8	9	9	10	11	12
56GB	1	2	3	3	4	5	6	6	7	8	9	9	10	11	12

 ©2024 Palo Alto Networks, Inc.

Memory Profile	Total vCPUs on Dataplane (17-32)															
	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
4.5GB	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
5.5GB	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
6.5GB	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
9GB	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
16GB	12	12	12	12	12	12	12	12	12	12	12	12	12	12	12	12
56GB	12	13	14	15	15	16	17	18	18	19	20	21	21	22	23	24

Continue to Software NGFW tasks:

- Maximum Limits Based on Memory
- Activate Credits
- Create a Deployment Profile
- Manage a Deployment Profile
- Register the VM-Series Firewall (Software NGFW Credits)
- Provision Panorama
- Migrate Panorama to a FW-Flex License
- Transfer Credits
- Renew Your Software NGFW Credit License
- Deactivate License (Software NGFW Credits)
- Create and Apply a Subscription-Only Auth Code
- Migrate to a Flexible VM-Series License

Maximum Limits Based on Memory

These limits apply to flexible licenses for VM-Series firewalls running PAN-OS 10.0.4 or later.

The following tables provide the maximum number for a particular object or resource that a single VM-Series firewall deployment can create, store, manage, or interact with on a firewall configured with 4.5, 5.5, 6.5, 9, 16, or 56 GB memory.

The memory profile and the total number of vCPUs determine how many cores are automatically assigned to the management plane and the dataplane.

If you are using SW NGFW licensing you can choose a memory profile that supports your requirements for one or more of the following resources:

• Address Assignment	• Interfaces	• Policies	• SSL Decryption

 ©2024 Palo Alto Networks, Inc.

- App-ID
- EDL
- GlobalProtect Client VPN
- GlobalProtect Clientless VPN
- High Availability
- IPSec VPN
- L2 Forwarding
- Multicast
- NAT
- Objects (addresses and services)
- QoS
- Routing
- Security Profiles
- Security Zones
- Sessions
- URL Filtering
- User-ID
- Virtual Routers
- Virtual Systems
- Virtual Wires

Sessions

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max sessions (IPv4 or IPv6)	50,000	64,000	250,000	819,200	2,000,000	10,000,000

Policies

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Security rules	200	250	1,500	10,000	10,000	20,000
Security rule schedules	256	256	256	256	256	256
NAT rules	400	400	3,000	5,000	8,000	15,000
Decryption rules	100	100	1,000	1,000	1,000	2,000
App override rules	100	100	1,000	1,000	1,000	2,000
Tunnel content inspection rules	100	100	100	500	500	2,000
SD-WAN rules	NA	100	100	100	300	300
Policy based forwarding rules	100	100	100	500	500	2,000
Captive portal rules	10	10	1,000	1,000	1,000	2,000
DoS protection rules	100	100	1,000	1,000	1,000	1,000

Security Zones

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max security zones	15	15	40	40	200	200

 ©2024 Palo Alto Networks, Inc.

Objects (addresses and services)

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Address objects	2,000	2,500	10,000	10,000	20,000	40,000
Address groups	100	125	1,000	1,000	2,500	4,000
Members per address group	2,500	2,500	2,500	2,500	2,500	2,500
Service objects	1,000	1,000	2,000	2,000	2,000	5,000
Service groups	250	250	500	500	250	500
Members per service group	500	500	500	500	500	500
FQDN address objects	2,000	2,000	2,000	2,000	2,000	2,000
Max DAG IP addresses* (system wide capacity)	1,000	1,000	2,500	200,000	300,000	300,500
Tags per IP address	32	32	32	32	32	32

* Firewall throughput measured with App-ID and User-ID features enabled utilizing AppMix transactions.

Security Profiles

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Security Profiles	38	38	375	375	750	750

App-ID

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Custom App-ID signatures	6,000	6,000	6,000	6,000	6,000	6,000
Shared custom App-IDs	512	512	512	512	512	512
Custom App-IDs (virtual system specific)	3,208	1,000	6,416	1,000	6,416	6,416

 ©2024 Palo Alto Networks, Inc.

User-ID

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
IP-User mappings (management plane)	524,288	524,288	524,288	524,288	524,288	524,288
IP-User mappings (data plane)	64,000	64,000	64,000	64,000	512,000	512,000
Active and unique groups used in policy (aggregate of LDAP groups, XML API Groups, and Dynamic User Group).*	1,000	1,000	1,000	1,000	10,000	10,000
Number of User-ID agents	100	100	100	100	100	100
Monitored servers for User-ID	100	100	100	100	100	100
Terminal server agents	400	400	400	400	2,000	2,500
Tags per User* (PAN-OS 9.1 and later)	32	32	32	32	32	32

*Firewall throughput measured with App-ID and User-ID features enabled utilizing AppMix transactions.

SSL Decryption

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max SSL inbound certificates	1,000	1,000	1,000	1,000	1,000	1,000
SSL certificate cache (forward proxy)	128	128	128	2,000	4,000	8,000
Max concurrent decryption sessions	1,024	1,024	6,400	15,000	50,000	100,000
SSL Port Mirror	Yes	Yes	Yes	Yes	Yes	Yes
SSL Decryption Broker	No	No	No	Yes	No	Yes
HSM Supported	Yes	Yes	Yes	Yes	Yes	Yes

 ©2024 Palo Alto Networks, Inc.

URL Filtering

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Total entries for allow list, block list and custom categories	25,000	25,000	25,000	25,000	25,000	100,000
Max custom categories	2,849	2,849	2,849	2,849	2,849	2,849
Max custom categories (virtual system specific)	500	500	500	500	500	500
Dataplane cache size for URL filtering	90,000	90,000	90,000	90,000	90,000	250,000
Management plane dynamic cache size	100,000	100,000	100,000	100,000	100,000	600,000

EDL

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max number of custom lists	30	30	30	30	30	30
Max number of IPs per system	50,000	50,000	50,000	50,000	50,000	50,000
Max number of DNS Domains per system	50,000	50,000	50,000	5000,000	2,000,000	2,000,00
Max number of URL per system	50,000	50,000	50,000	100,000	100,000	100,000
Shortest check interval (min)	5	5	5	5	5	5

Interfaces

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Mgmt - out-of-band	NA	NA	NA	NA	NA	NA
Mgmt - 10/100/1000 high availability	NA	NA	NA	NA	NA	NA
Mgmt - 40Gbps high availability	NA	NA	NA	NA	NA	NA

 ©2024 Palo Alto Networks, Inc.

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Mgmt - 10Gbps high availability	NA	NA	NA	NA	NA	NA
Traffic - 10/100/1000	NA	NA	NA	NA	NA	NA
Traffic - 100/1000/10000	NA	NA	NA	NA	NA	NA
Traffic - 1Gbps SFP	NA	NA	NA	NA	NA	NA
Traffic - 10Gbps SFP+	NA	NA	NA	NA	NA	NA
Traffic - 40/100Gbps QSFP+/ QSFP28	NA	NA	NA	NA	NA	NA
802.1q tags per device	4,094	4,094	4,094	4,094	4,094	4,094
802.1q tags per physical interface	4,094	4,094	4,094	4,094	4,094	4,094
Max interfaces (logical and physical)	512	512	2,048	2,048	4,096	40,96
Maximum aggregate interfaces	NA	NA	NA	NA	NA	NA
Maximum SD-WAN virtual interfaces	NA	150	300	500	1,000	1,000

Virtual Routers

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Virtual routers	3	3	3	10	20	125

Virtual Wires

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Virtual wires	2	4	12	12	12	12

Virtual Systems

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Base virtual systems	1	1	1	1	1	1
Max virtual systems	NA	NA	NA	NA	NA	NA

©2024 Palo Alto Networks, Inc.

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Additional licenses are required for virtual system capacities above the base virtual system's capacity						

Routing

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
IPv4 forwarding table size* (Entries shared across virtual routers)	1,000	2,500	5,000	10,000	32,000	100,000
IPv6 forwarding table size* (Entries shared across virtual routers)	1,000	1,000	5,000	10,000	32,000	100,000
System total forwarding table size	1,000	1,000	5,000	10,000	32,000	100,000
Max route maps per virtual router	50	50	50	50	50	50
Max routing peers (protocol dependent)	500	500	500	500	1,000	1,000
Static entries - DNS proxy	1,024	1,024	1,024	1,024	1,024	1,024
Bidirectional Forwarding Detection (BFD) Sessions	N/A	N/A	128	512	1,024	1,024

*Firewall throughput measured with App-ID and User-ID features enabled utilizing AppMix transactions.

L2 Forwarding

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
ARP table size per device	1,500	1,500	2,500	10,000	32,000	128,000
IPv6 neighbor table size	500	500	2,500	10,000	32,000	128,000
MAC table size per device	1,500	1,500	2,500	5,000	32,000	128,000
Max ARP entries per broadcast domain	1,500	1,500	2,500	10,000	32,000	128,000

 ©2024 Palo Alto Networks, Inc.

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max MAC entries per broadcast domain	1,500	1,500	2,500	5,000	32,000	128,000

NAT

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Total NAT rule capacity	160	400	3,000	5,000	8,000	8,000
Max NAT rules (static)* (Configuring static NAT rules to full capacity requires that no other NAT rule types are used.)	160	400	3,000	5,000	8,000	8,000
Max NAT rules (DIP)* (Configuring DIP NAT rules to full capacity requires that no other NAT rule types are used.)	160	400	2,000	3,000	8,000	8,000
Max NAT rules (DIPP)	160	200	400	800	2,000	2,000
Max translated IPs (DIP)	16,000	16,000	128,000	128,000	160,000	160,000
Max translated IPs (DIPP)* (DIPP translated IP capacity is proportional to the DIPP pool oversubscription value. The capacity shown here is based on an oversubscription value of 1x.)	200	200	400	800	2,000	2,000
Default DIPP pool oversubscription* (Source IP and source port reuse across concurrent sessions)	2	2	2	2	8	8

*Firewall throughput measured with App-ID and User-ID features enabled utilizing AppMix transactions.

Address Assignment

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
DHCP servers	3	3	3	10	20	125

 ©2024 Palo Alto Networks, Inc.

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
DHCP relays* (Maximum capacity represents total DHCP servers and DHCP relays combined)	500	500	500	500	500	500
Max number of assigned addresses	64,000	64,000	64,000	64,000	64,000	64,000

*Firewall throughput measured with App-ID and User-ID features enabled utilizing AppMix transactions.

High Availability

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Devices supported	2	2	2	2	2	2
Max virtual addresses	32	32	128	32	32	128

QoS

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Number of QoS policies	100	100	500	1,000	2,000	4,000
Physical interfaces supporting QoS	6	6	6	6	12	12
Clear text nodes per physical interface	31	31	31	63	63	63
DSCP marking by policy	Yes	Yes	Yes	Yes	Yes	Yes
Subinterfaces supported	NA	NA	NA	NA	NA	NA

IPSec VPN

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max IKE Peers	25	250	1,000	1,000	1,000	2,000
Site to site (with proxy id)	25	250	1,000	2,000	4,000	8,000
SD-WAN IPSec tunnels	NA	250	1,000	1,000	1,000	2,000

©2024 Palo Alto Networks, Inc.

GlobalProtect Client VPN

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max tunnels (SSL, IPSec, and IKE with XAUTH)	25	250	500	2,000	6,000	12,000

GlobalProtect Clientless VPN

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Max SSL tunnels	40	40	100	400	1,200	2,500

Multicast

Feature	4.5 GB	5.5 GB	6.5 GB	9 GB	16 GB	56 GB
Replication (egress interfaces)	100	100	100	100	100	100
Routes	500	500	2,000	2,000	4,000	4,000

Activate Credits

Within your organization you can create many accounts, each with a different purpose. During activation you can choose only one account per default credit pool. Once you the credit pool is active, users granted the credit administrator role can allocate the credits for deployments, and even transfer credits to other pools.

If you have an existing CSP account and are a superuser or an admin, the system automatically adds the credit admin role to your profile. If you do not have an existing account, the CSP automatically creates an account for you and adds the credit admin role to your profile.

You (the purchaser) receive an email detailing the subscription, the credit pool ID, the subscription start and end date, the amount of credits purchased, and the description of the default credit pool (see “default credit pools” in VM-Series Firewall Licensing).

Secure this email for future reference.

STEP 1 | In the email, click **Start Activation** to view your available credit pools.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Select the credit pool you want to activate. You can use the search field to filter your account list by number or name.

If you have purchased multiple credit pools (see Software NGFW Credits), they are automatically selected. The check marks represent activation links for onboarding credits.

You are prompted to authenticate or sign in.

*If you deselect a credit pool, you see a reminder that if you want to activate those credits, you must return to the email and click the **Start Activation** link.*

STEP 3 | Select **Start Activation**.

STEP 4 | Select the support account (you can search by account number or name).

STEP 5 | Select the default credit pool.

STEP 6 | Select **Deposit Credits**.

You see a message that the deposit was successful.

STEP 7 | (optional) If this is your first credit activation, you see the **Create Deployment Profile** dialog.

Continue to Create a Deployment Profile.

Create a Deployment Profile

To create a deployment profile, you must have a Customer Support Portal account (CSP) and access to an activated credit pool.

Before you begin, estimate the number of firewalls that will use the configuration in the deployment profile. You don't have to deploy all the firewalls at once.

STEP 1 | If you already have a credit pool, log into the account, and from the dashboard, select **Assets** > **Software NGFW Credits** > **Create Deployment Profile**.

If you have just activated a credit pool you see the **Create Deployment Profile** form.

1. Select the VM-Series firewall type.
2. Select the PAN-OS version.
   - **Fixed Models** (VM-Series Models)
   - **Flexible vCPUs (PAN-OS 10.0.4 and above)**
3. Click **Next**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | VM-Series profile.

1. **Profile Name**.

   Name the profile.

2. **Number of Firewalls**.

   Enter the number of firewalls this profile deploys, assuming you have sufficient credits. You do not have to deploy them all at once.

3. **Firewall Model**:

   Choose a VM-Series model.

   Planned vCPU/Firewall (PAN-OS 10.0.4 or above).

   Enter the number of vCPUs per firewall.

   Security Use Case: Choose a use case.

4. **Customize Subscriptions**.

   After selecting a use case, you can add or remove security services.

5. (optional) **Use Credits to Enable VM Panorama**.

   Choose the Panorama use case(s)—Management and/or Log Collector.

STEP 3 | (optional) Hover over the question mark following **Protect more, save more** to see how your credit allocation affects savings.

STEP 4 | Click **Calculate Estimated Cost** to view the credit total, and the number of credits available before the deployment.

(optional) Hover over the question mark following the estimate to view the credit breakdown for each component.

STEP 5 | **Create the Deployment Profile**.

You might have to wait several seconds for the profile to appear in the **Current Deployment Profiles** tab list. Before the allocation is complete, the **Credits Consumed/Allocated** column shows 0 and **Update Pending**.Scroll to the bottom and go to the last page to find your profile.

To view your deployment profile later on, click the **Details** button on the parent credit pool and select **Current Deployment Profiles**.

- Note the Auth Code for your profile on the far right; Software NGFW credit auth codes start with D.
- The **Credits Consumed/Allocated** column shows 0 and **Update Pending** before the allocation is complete.
- The **Audit Trail** tab shows **Credit Transactions** and the **Deployment Profiles** you manage. You can also search for a profile by time in this tab.

  Use search to locate your profile, and expand the row to view the configuration you specified when you created the profile.

Manage a Deployment Profile

After you create your deployment profile you can edit, copy, or delete it.

 ©2024 Palo Alto Networks, Inc.

- Edit a Deployment Profile
- Clone a Deployment Profile
- Transfer a Deployment Profile
- Delete a Deployment Profile

Edit a Deployment Profile

STEP 1 | Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your profile.

STEP 2 | Select the **Current Deployment Profiles** tab.

STEP 3 | On the far right, select the vertical ellipsis (More Options) and select **Edit Profile**.

STEP 4 | Make your changes and select **Update Deployment Profile**.

STEP 5 | Select the **Audit Trail** tab and use search to locate your profile.

Use search to locate your profile, and expand the row to view the configuration you specified when you created the profile.

Clone a Deployment Profile

STEP 1 | Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your profile.

STEP 2 | On the far right, select the vertical ellipsis (More Options) and select **Clone Profile**.

STEP 3 | Change the profile name, make any other changes, and select **Create Deployment Profile**.

STEP 4 | Select the **Audit Trail** tab and use search to locate your profile.

Expand the row to view the configuration you cloned. It is a new configuration with a different Profile Name and auth code.

Transfer a Deployment Profile

Use the following procedure to transfer a deployment profile from one credit pool to another.

STEP 1 | Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your profile.

STEP 2 | On the far right, select the vertical ellipsis (More Options) and select **Transfer Profile**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Select the target credit pool and click **Transfer**.

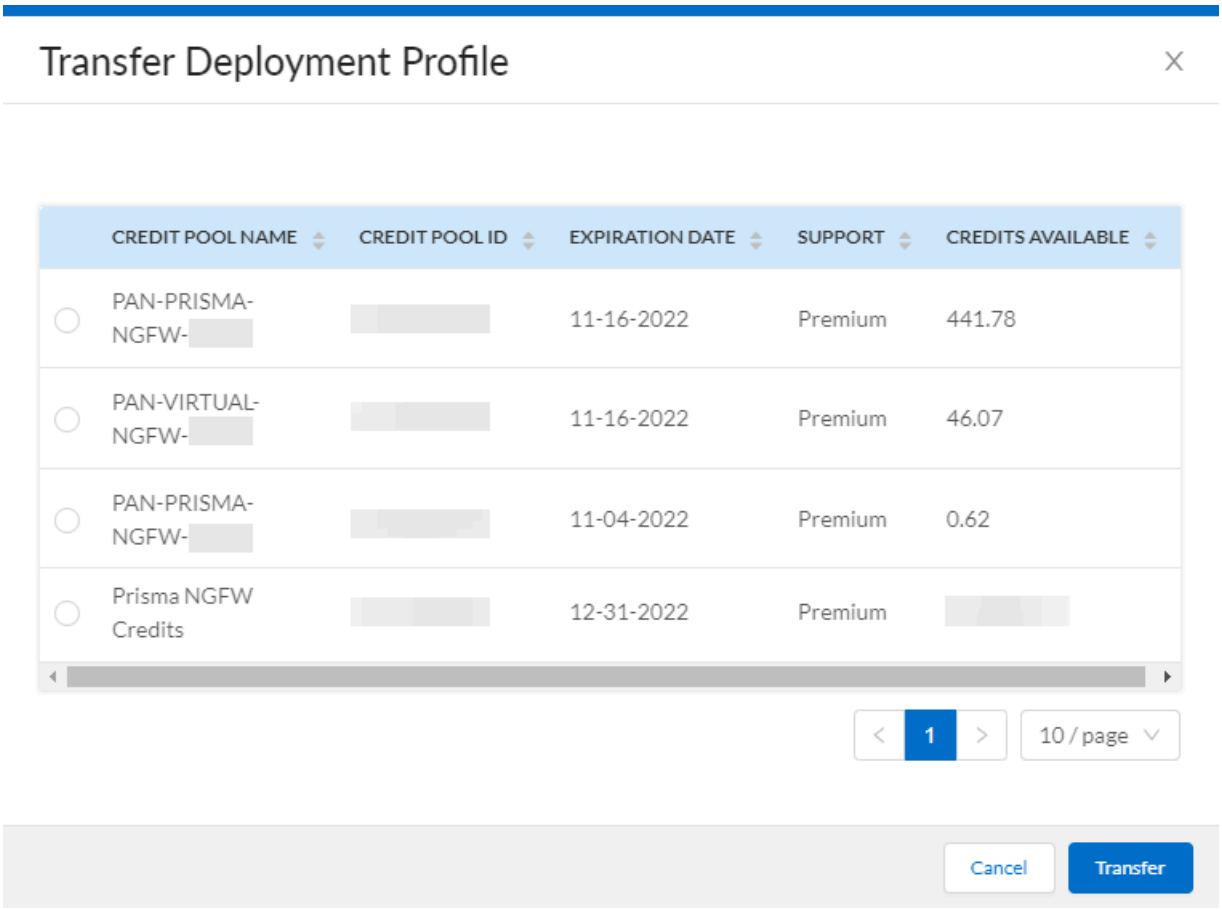

Delete a Deployment Profile

Before deleting a deployment profile, you must Deactivate License (Software NGFW Credits) on any firewall using the deployment profile and then deactivate the VM.

STEP 1 | Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your profile.

STEP 2 | On the far right, select the vertical ellipsis (More Options) and select **Delete**.

Register the VM-Series Firewall (Software NGFW Credits)

Registration requires access to the Palo Alto Networks customer support portal (the CSP) and a support account. Create one if necessary.

During activation, an administrator activates a credit pool and the credits are deposited. When anyone creates a deployment profile, an auth code is created. Complete one of the following procedures to initiate registration.

- Device can access CSP
- Device cannot access the CSP

 ©2024 Palo Alto Networks, Inc.

- **Use the following steps if the firewall is able to connect to the CSP:**
  1. Log in to the CSP with your account credentials.
  2. Select **Assets** > **Software NFGW Credits**.

     Locate your credit pool and view **Details**.
  3. View **Current Deployment Profiles** and choose (or create) a profile.

     You will use the auth code from this profile for licensing any firewall you create with it. An auth code for a flexible firewall license begins with the letter D.
  4. Log in to the VM-Series firewall web interface.
  5. Verify the Palo Alto Networks update server configuration.
     1. Select **Device** > **Setup** > **Services**.
     2. Confirm that **Update Server** is set to `updates.paloaltonetworks.com`.
     3. Confirm that **Verify Update Server Identity** is selected.
  6. Select **Device** > **Licenses**.
     1. Select the **Activate feature using authorization code** link.
     2. Enter the VM-Series authorization code from the deployment profile.
     3. Click **OK** to confirm the license upgrade.The firewall contacts the Palo Alto Networks update server and consumes the tokens required for your firewall based on the VM-Series model.
  7. Confirm that the **Dashboard** displays a valid serial number and that the **PA-VM** license displays in the **Device** > **Licenses** tab.
  8. Verify your firewall is registered on the CSP:
     - Select **Assets** > **Software NFGW Credits**
     - Auth Code column, **View Devices** and locate the serial number for your deployment.
     - In the credit pool, Credits Consumed, Firewalls Deployed, and vCPUs consumed should be incremented to reflect your deployment.

 ©2024 Palo Alto Networks, Inc.

- **Use the following steps if the firewall is *not* able to connect to the CSP**:

  This workflow adds your firewall to the support database. Because the firewall can't connect to the license server, you must manually pass the licenses from the CSP to the firewall.

  1. Log in to the CSP with your account credentials.
  2. Select the new profile and Select the vertical ellipsis (More Options) and **Register Firewall**.

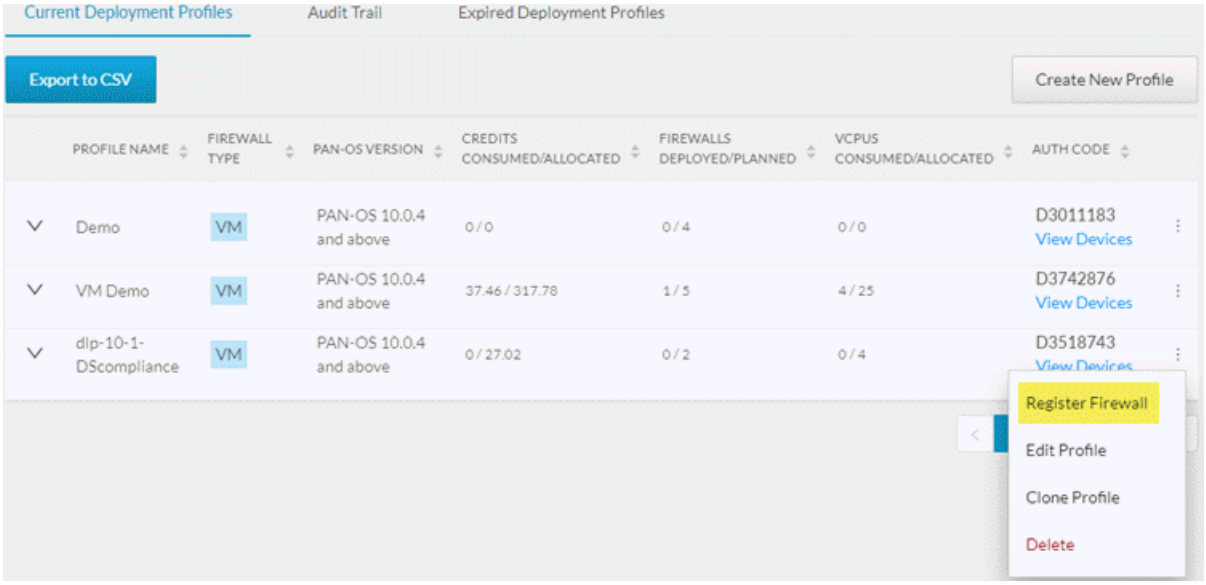

  This opens the device registration form. Enter the information for your firewall and **Submit**:

 ©2024 Palo Alto Networks, Inc.

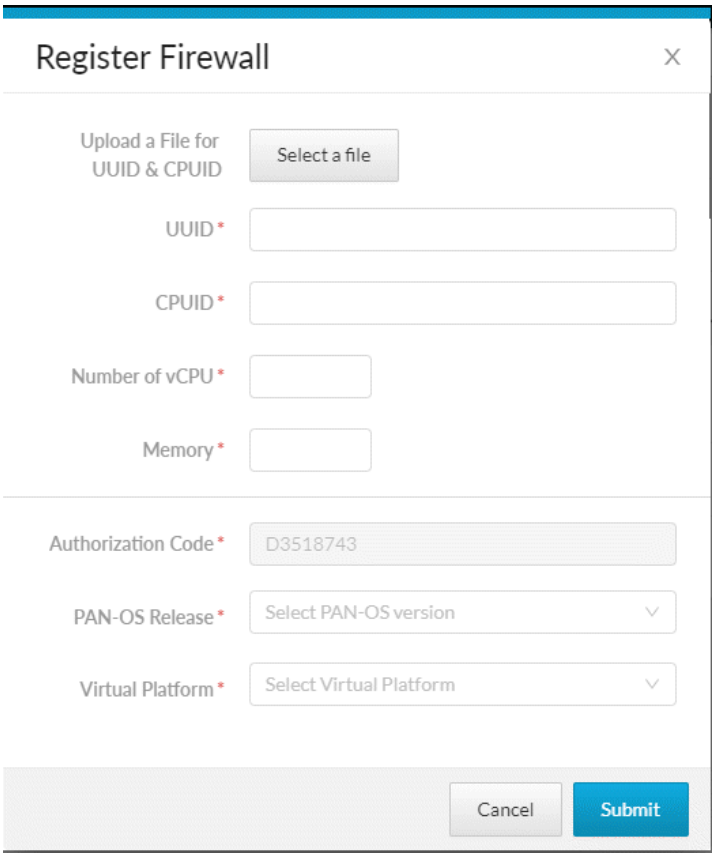

This associates the firewall with the profile and its authcode and assigns a serial number.

3. Click **View Devices** to see associated firewalls in **Software NGFW Devices**.

   In the **License** column, download each license key to a location from which you can safely transfer the files to the firewall.

4. Log in to the firewall and select **Device** > **Licenses**.

   License keys must be installed through the web interface. The firewall does not support license key installation through SCP or FTP.

   - Click **Manually Upload License**.
   - Confirm that the Dashboard displays a valid serial number and that the **PA-VM** license displays in the **Device** > **Licenses** tab.

Provision Panorama

This option is only visible if you selected Panorama when the deployment profile was created. You can edit the profile, if necessary.

STEP 1 | Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your profile.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | On the far right, select the vertical ellipsis (More Options) and select **Provision Panorama**, and **Provision**. You see the list of firewalls provisioned for the current deployment profile.

This creates a Panorama, assigns a serial number and the model type PAN-PRA-1000-CP, and registers the Panorama as an asset. The Panorama you just provisioned is the last Panorama listed. Note the auth code starts with F (so it is not the same as the deployment profile), but the expiration date is the same as your profile's credit pool.

Copy the serial number.

STEP 3 | From the deployment profile, **View Devices**, and select **Panorama** on the Software NGFW Devices page. This displays all SW NGFW Panoramas.

Search By serial number using the serial number you copied.

You can also select **Assets** > **Software NGFW Devices** and **Search By** serial number with the serial number you copied.

STEP 4 | After setting up your Panorama Virtual Appliance, add the serial number to Panorama.

1. Log in to Panorama.
2. Select **Panorama** > **Setup** > **Management** > **General Settings** and click the **Edit** icon.
3. Enter the serial number you copied from the CSP in the **Serial Number** field.
4. Click **OK** to save your changes.
5. Commit your configuration changes.

   Select **Commit** > **Commit to Panorama** and **Commit** your changes.

Migrate Panorama to a FW-Flex License

Complete the following procedure to migrate your VM-ELA or perpetual virtual Panorama license to a FW-Flex license. This migration allows you to move your existing Panorama devices to the FW-Flex license without disruption while retaining your existing serial number. Because your serial number does not change, your logs and existing policy are retained.

STEP 1 | Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your profile.

STEP 2 | On the far right, select the vertical ellipsis (More Options) and select **Provision Panorama** and then click **Migrate Existing**.

The CSP displays all virtual Panorama devices associated with your support account.

STEP 3 | Select the check box for each virtual Panorama to be migrated.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Click **Migrate**.

Verify that the **Current Support Expiration Date** has been updated. Additionally, you can expand each line to view the individual licenses applied to the selected Panorama.

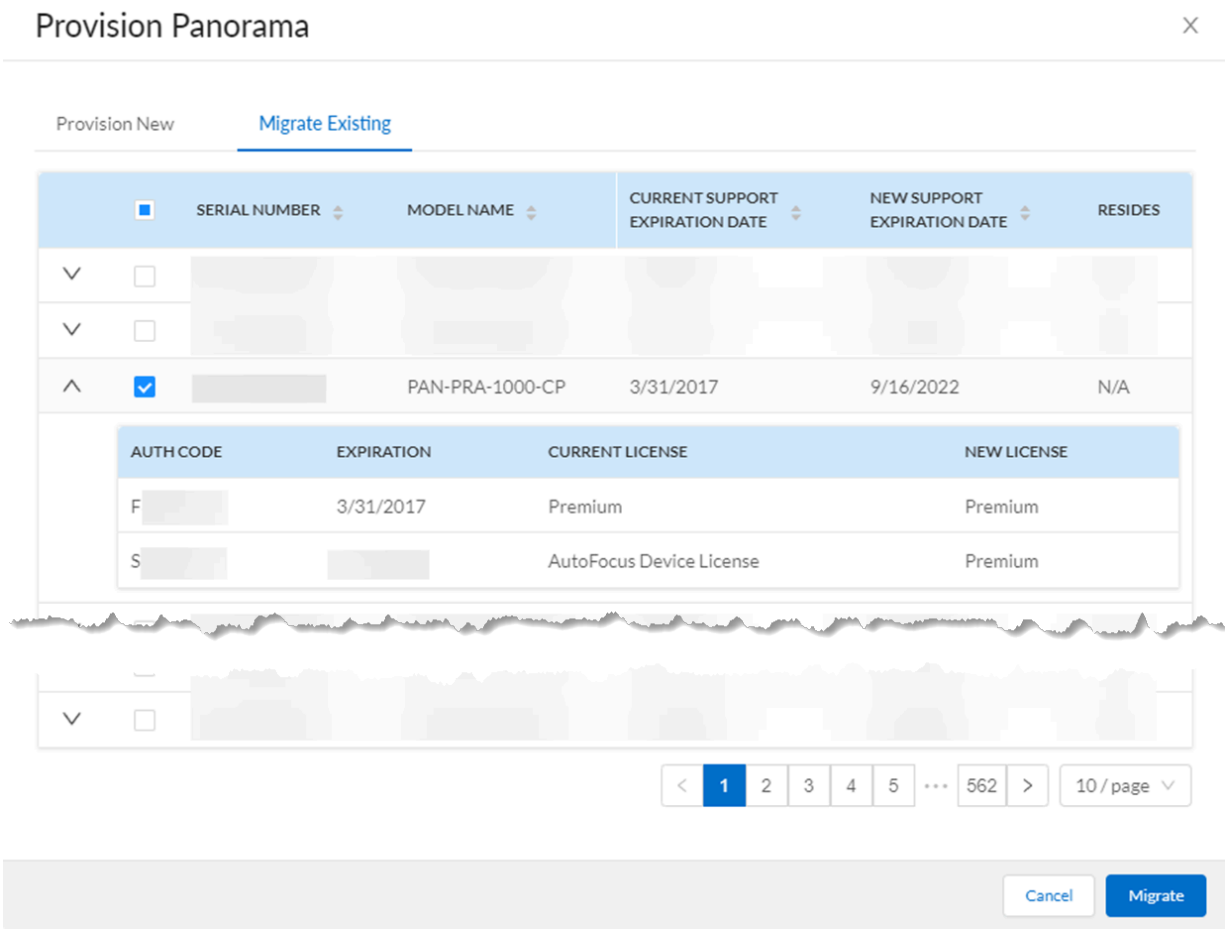

Transfer Credits

From the Customer Support Portal (CSP), transfer credits to a credit pool in the same account, or to a credit pool in a different account that you can access.

Credits must be transferred between the pools within the same contract (parent/child).

- Different CSP Account
- Different Pool in this Account

 ©2024 Palo Alto Networks, Inc.

Different CSP Account

STEP 1 | Log into your CSP account.

STEP 2 | Select **Assets** > **Software NGFW Credits**.

- Identify the source credit pool and make note of the Credit Pool ID.
- Identify the destination credit pool and make note of the Credit Pool ID.

  If the destination is in a different account, select it from the **Current Account** dropdown on the upper left, and Select **Assets** > **Software NGFW Credits**. Find the destination and note the credit type and the Credit Pool ID.

STEP 3 | Go to the source credit pool and click **Transfer Credits** on the bottom left.

STEP 4 | Choose Different CSP account.

1. **Transfer to**—Choose an account name.
2. **As credit type**—Choose a credit type. At this time, the source and destination type must be the same.
3. **Credit Pool ID#**—Choose a Credit pool ID number.

   If the destination account does not have any credit pools of the chosen type, the CSP prompts you to create a credit pool.
4. **Amount to transfer**—Enter the amount to transfer.

STEP 5 | Select **Update Credits**.

You might need to wait a short time or refresh your screen to see the change.

STEP 6 | To view credit transactions for a pool, select **Details** and select **Audit Trail**.

Different Pool in this Account

STEP 1 | Log in to your CSP account.

STEP 2 | Select **Assets** > **Software NGFW Credits**.

- Identify the destination credit pool and make note of the Credit Pool ID.
- If there isn't a destination credit pool of the type you specify, you are prompted to create a new credit pool.

STEP 3 | Go to the source credit pool and select **Transfer Credits** on the bottom left.

STEP 4 | Select **Different Pool in this Account**.

1. **New credit type**—Choose a credit type. At this time, the source and destination type must be the same.
2. **Credit Pool ID#**—Choose a Credit pool ID number.

   If the destination account does not have any credit pools of the chosen type, the CSP prompts you to create a credit pool.
3. **Amount to transfer**—Enter the amount to transfer.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Select **Update Credits**.

You might need to wait a short time or refresh your screen to see the change.

STEP 6 | To view credit transactions for a pool, select **Details** and select **Audit Trail**.

If you want to transfer credits between pools, the expiration dates must be the same on both credit pools.

Renew Your Software NGFW Credit License

When a deployment profile expires, it moves to the from the Current Deployment Profiles tab to the Renew Profiles tab. From the Renew Profiles tab, you can quickly renew your deployment profiles from the Palo Alto Networks Customer Support Portal without disrupting operations of your VM-Series firewall. After a deployment profile expires and moves to the Renew Profiles tab, you have 30 days to renew the profile. Any deployment profiles not renewed within 30 days expire and move to the Expired Deployment Profile tab. See What Happens When Licenses Expire? for more information.

After renewal, you might notice some changes between your deployment profiles before. The Prisma NGFW Credits and Virtual NGFW Credits credit pools are now both called Software NGFW Credits credit pools. Additionally, the number of credits in your pool might change after renewal due to changes in the product pricing model.

STEP 1 | Log in to the Palo Alto Networks Customer Support Portal.

STEP 2 | Select **Assests** > **Software NGFW Credits**.

STEP 3 | Locate the Deployment Profile to be renewed and click **Details**.

STEP 4 | Select **Renew Profiles**.

STEP 5 | Click the Renew icon and then click **Renew** to confirm.

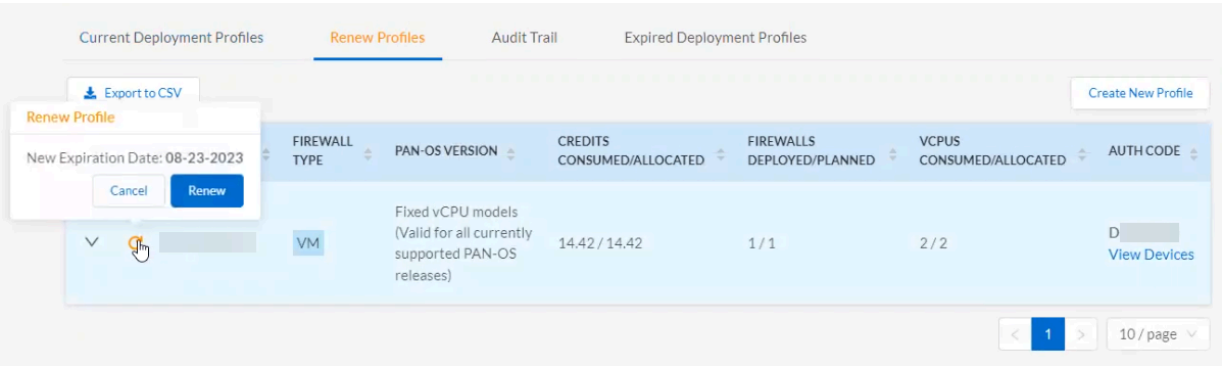

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Verify that your deployment profile was renewed successfully.

1. Click **Current Deployment Profiles**.
2. Confirm that the renewed deployment profile is displayed.

   Additionally, you can return to the Software NGFW dashboard to view your credit pool. After a successful renewal, the credit pool displays **Renewal Confirmed**. This message remains until the end of the expiry date of your renewed Software NGFW credits.

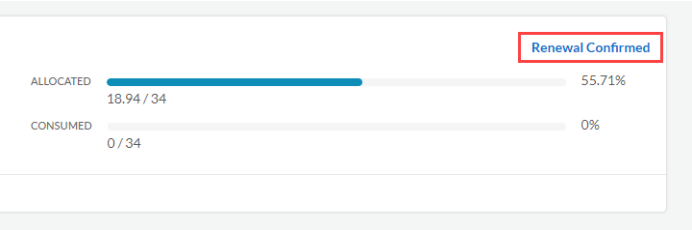

Deactivate License (Software NGFW Credits)

You must deactivate any licenses from the CSP **before** you delete a firewall (or the VM hosting the firewall) or the license credits cannot return to your credit pool.

When you have internet access to the licensing server, deactivating the firewall on the CSP automatically removes the licenses and the remaining credits are returned to the deployment profile. After you deactivate the license you must delete the firewall or it will continue to consume credits.

If you don't have internet access, you must export the license token from the firewall. Then, in the CSP, start deactivation and upload the token (or paste in the token text) to complete the deactivation.

- Internet access
- No internet access
- No internet access - Panorama management

- Direct internet access.
  1. Select **Assets** > **Software NGFW Credits** and click the **Details** button on the credit pool you used to create your deployment profile.
  2. Locate your deployment profile, and on the far right, select the vertical ellipsis (More Options) and select **Deactivate Firewall**.
  3. Check the firewall you want to deactivate, and select **Deactivate Firewall**.

 ©2024 Palo Alto Networks, Inc.

- No internet access.
  1. Log in to the firewall web interface and select **Device** > **Licenses**.
  2. In the License Management section, select **Deactivate VM**.

     Verify the list of licenses/entitlements to be deactivated on the firewall.
  3. Select **Complete Manually** to start the deactivation.

     Click the **Export license token** link to save the token file to the client. A token filename looks like this: `20150128_1307_dact_lic.01282015.130737.tok`

     At this point the license has been deactivated on the firewall, but the credits have not been returned to the credit pool.
  4. Use the token file to register the changes with the Licensing server:

     1. Log into the Palo Alto Networks Customer Support website.
     2. Select **Assets** > **VM-Series Auth-Codes** > **Deactivate License(s)**.

     In the Deactivate Licenses form, paste in the token text, or copy the token to a computer with internet access and upload the token file to the CSP to complete license removal.
  5. Delete the VM

- No internet access—Panorama management
  1. Log in to the Panorama web interface and select **Panorama** > **Device Deployment** > **Licenses**.
  2. **Select Deactivate VMs** and select the VM-Series firewall that you want to deactivate.
  3. Select **Complete Manually** to export the token file.
  4. Click the **Export license token** link to save the token file. A token filename looks like this: 20150128_1307_dact_lic.01282015.130737.tok

     If the export is successful, a completion message is displayed, and the firewall reboots automatically.
  5. Use the token file to register the changes with the licensing server.

     1. Log into the Palo Alto Networks Customer Support website.
     2. Select **Assets** > **VM-Series Auth-Codes** > **Deactivate License(s)**.

     In the Deactivate Licenses form, paste in the token text, or copy the token to a computer with internet access and upload the token file to the CSP to complete license removal.
  6. (optional) Remove the deactivated VM-Series firewall as a managed device on Panorama.

     Instead of deleting the deactivated firewalls, you can create a separate device group and assign them to it.

     1. Select **Panorama** > **Managed Devices**.
     2. Select the firewall that you deactivated and click **Delete**.

 ©2024 Palo Alto Networks, Inc.

Create and Apply a Subscription-Only Auth Code

If you have already deployed VM-Series firewalls using a perpetual or ELA license bundle, you can use Software NFGW credits to add subscriptions to your existing deployment.

When adding subscriptions, keep the following in mind:

- Only select subscriptions that are not included in your existing bundle.
- Ensure that your VM-Series model and PAN-OS version support the subscription you want to add.
- Do not apply more than one subscription-only auth code to a firewall.
- A subscription-only license cannot be used as a renewal for an existing subscription.
- The firewall on which you are applying the auth code must be able to access the Palo Alto Networks licensing server.

STEP 1 | If you already have a credit pool, log into the account, and from the dashboard, select **Assets** > **Software NGFW Credits** > **Create Deployment Profile**.

If you have just activated a credit pool you see the **Create Deployment Profile** form.

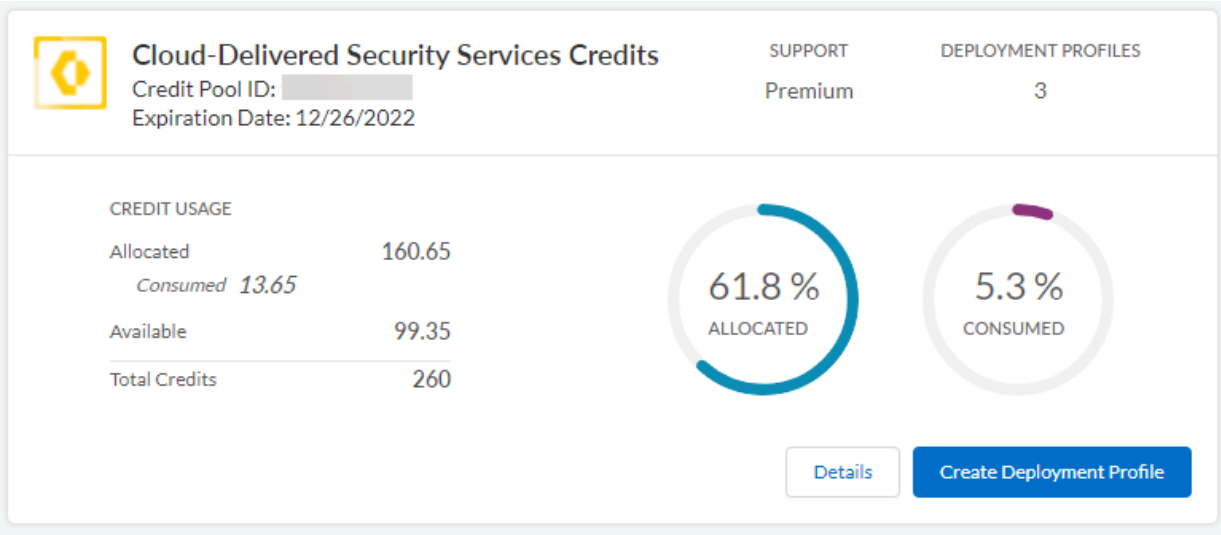

STEP 2 | Click **Next**. A subscription-only deployment profile can only be used on VM-Series firewalls with

 ©2024 Palo Alto Networks, Inc.

STEP 3 | VM-Series profile.

1. **Profile Name**.

   Name the profile.

2. **Number of Firewalls**.

   Enter the number of firewalls this profile deploys, assuming you have sufficient credits. You do not have to deploy them all at once.

3. **Firewall Model**:

   Choose the VM-Series model of the VM-Series firewalls to which you are adding subscriptions.

   You can only apply the subscription-only auth code to VM-Series firewall of the chosen model.

4. **Security Use Case:** Choose a use case.

5. **Customize Subscriptions**.

   After selecting a use case, you can add or remove security services.

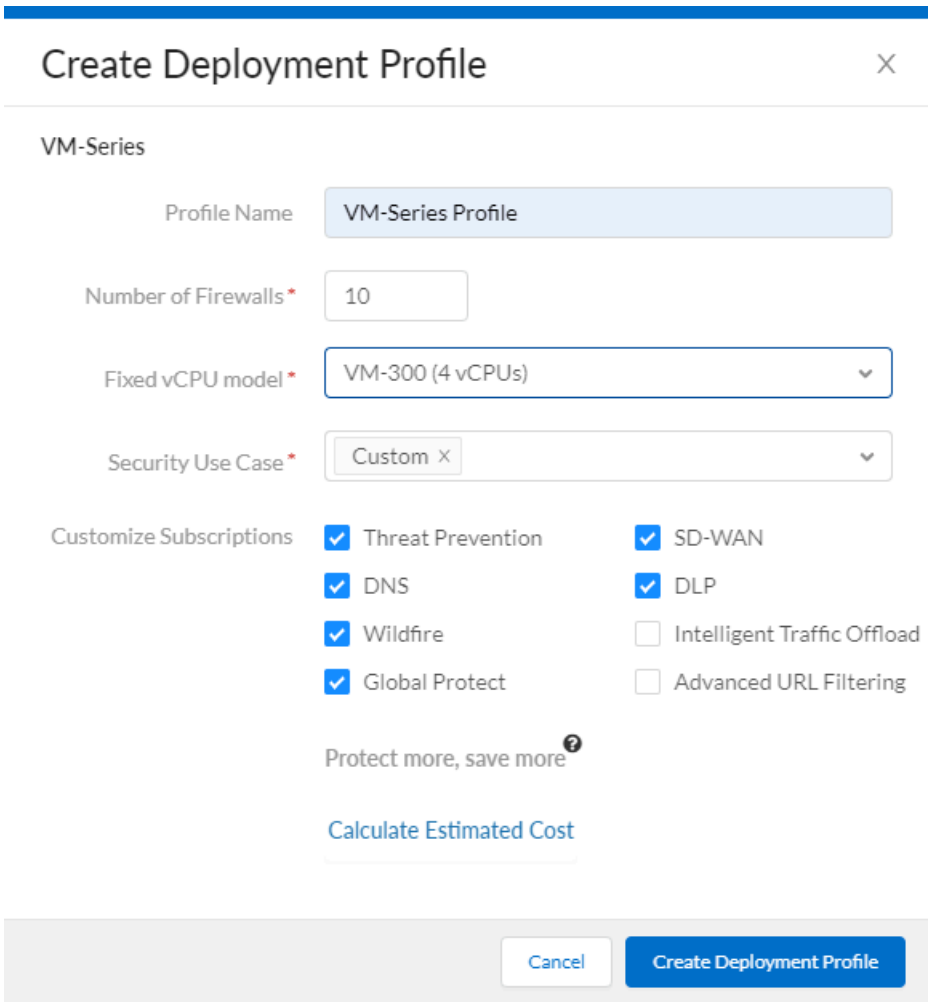

 ©2024 Palo Alto Networks, Inc.

STEP 4 | (optional) Hover over the question mark following **Protect more, save more** to see how your credit allocation affects savings.

STEP 5 | Click **Calculate Estimated Cost** to view the credit total, and the number of credits available before the deployment.

(optional) Hover over the question mark following the estimate to view the credit breakdown for each component.

STEP 6 | **Create the Deployment Profile**.

You might have to wait several seconds for the profile to appear in the **Current Deployment Profiles** tab list. Before the allocation is complete, the **Credits Consumed/Allocated** column shows 0 and **Update Pending**. Scroll to the bottom and go to the last page to find your profile.

To view your deployment profile later on, click the **Details** button on the parent credit pool and select **Current Deployment Profiles**.

- Note the Auth Code for your profile on the far right; subscription-only auth codes start with D.
- The **Credits Consumed/Allocated** column shows 0 and **Update Pending** before the allocation is complete.
- The **Audit Trail** tab shows **Credit Transactions** and the **Deployment Profiles** you manage. You can also search for a profile by time in this tab.

  Use search to locate your profile, and expand the row to view the configuration you specified when you created the profile.

STEP 7 | Apply the subscription-only auth code to your firewall.

1. Log in to the CSP with your account credentials.
2. Select **Assets** > **Software NFGW Credits**.

   Locate your credit pool and view **Details**.
3. View **Current Deployment Profiles** and choose a profile.

   You will use the auth code from this profile on any firewall of the model specified when you created your deployment profile. An auth code for a subscription-only license begins with the letter D.
4. Log in to the VM-Series firewall web interface.
5. Verify the Palo Alto Networks update server configuration.
   1. Select **Device** > **Setup** > **Services**.
   2. Confirm that **Update Server** is set to `updates.paloaltonetworks.com`.
   3. Confirm that **Verify Update Server Identity** is selected.
6. Select **Device** > **Licenses**.
   1. Select the **Activate feature using authorization code** link.
   2. Enter the subscription-only auth code from the deployment profile.
   3. Click **OK** to confirm the license application.The firewall contacts the Palo Alto Networks update server and consumes the tokens required for your subscription-only license.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Verify that the license is successfully activated.

On the **Device** > **Licenses** page, verify that the license is successfully activated. For example, after activating the Threat Prevention license, you should see that the license is valid:

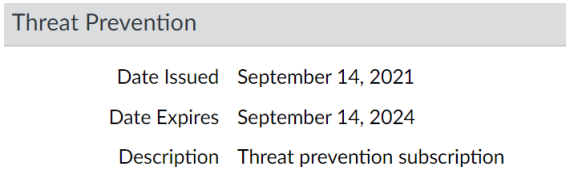

Migrate to a Flexible VM-Series License

You can migrate your VM-Series firewall perpetual or ELA license to a flexible VM-Series firewall license (funded using Software NGFW credits). You can switch the license on an individual firewall or on multiple firewalls simultaneously from Panorama.

*If the license you want to switch is for a VM-Series firewall with and ELA or perpetual license, you must choose **Fixed Models** and use the same VM-Series model and number of vCPUs when you create the SW NGFW deployment profile for the flexible license. For example, if you are currently using a VM-100 with 2 vCPUs, when you create the deployment profile, select **Fixed Models**, select VM-100, and specify 2 vCPUs. As long as the capacity is the same, you can add new services.*

Complete one of the following procedures to migrate your licenses.

- Firewall can access the CSP
- Firewall cannot access the CSP
- Panorama can access the CSP
- Panorama cannot access the CSP
- Verify the migration

- **Firewall can access the CSP**—Migrate a license on a standalone firewall

  This process does not disrupt traffic moving through the firewall.

  1. Log in to the VM-Series firewall web interface.
  2. Verify the Palo Alto Networks update server configuration.
     1. Select **Device** > **Setup** > **Services**.
     2. Confirm that **Update Server** is set to `updates.paloaltonetworks.com`.
     3. Confirm that **Verify Update Server Identity** is selected.
  3. Log in to the CSP and Create a Deployment Profile. If the license you want to migrate is for a VM-Series firewall with an ELA or perpetual license, you must choose **Fixed Models**

 ©2024 Palo Alto Networks, Inc.

and use the same VM-Series model and number of vCPUs when you create the SW NGFW deployment profile for the flexible license.

You will use the auth code from this profile. An auth code for a flexible firewall license begins with the letter D, as shown below.

 ©2024 Palo Alto Networks, Inc.

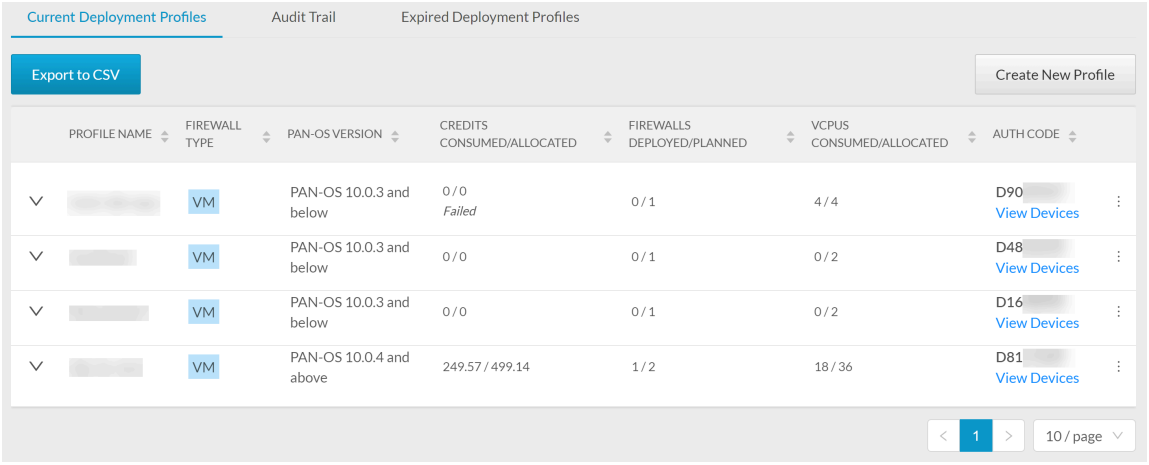

 ©2024 Palo Alto Networks, Inc.

1. Select **Device** > **Licenses**.

   If the current VM-Series model and the VM-Series model you are migrating to are different, select the **Upgrade VM Capacity** link.

   If the VM-Series model is the same before and after migration, select the **Activate feature using authorization code** link.
2. Enter the VM-Series authorization code from the new deployment profile.
3. Click **OK** to confirm the license upgrade.The firewall contacts the Palo Alto Networks update server and consumes the tokens required for your firewall based on the VM-Series model.
4. (Optional) Verify the migration.

4. Repeat this process for each VM-Series firewall in your deployment.

- **Firewall cannot access the CSP**—Migrate a license on an offline firewall

1. If necessary, install the license API key on your VM-Series firewall.
2. Use the CLI to use manual mode to deactivate the fixed-model license.
3. Deactivate the VM from the Firewall using the manual procedure, and log in to the CSP and use the token file to deactivate the VM.
4. In the CSP, Create a Deployment Profile with the same VM-Series model, number of vCPUs, and security subscriptions as the old fixed model license. You will use the auth code from this profile.
5. Select the new profile and click the vertical ellipsis and select **Register Firewall**.
   1. Enter the VM and firewall information and select **Submit**. This associates the firewall with the profile and its authcode and assigns a serial number.
   2. Click **View Devices** to see associated devices in **Software NGFW Devices**.
   3. In the **License** column, download the license keys to a location from which you can safely transfer the files to the host machine.
6. On the firewall select **Device** > **Licenses**.

   License keys must be installed through the web interface. The firewall does not support license key installation through SCP or FTP.
7. Click **Manually Upload License** and enter the license keys.
8. Confirm that the **Dashboard** displays a valid serial number and that the **PA-VM** license displays in the **Device** > **Licenses** tab.
9. (optional) Verify the migration.

 ©2024 Palo Alto Networks, Inc.

- **Panorama can access the CSP**—Migrate licenses on Panorama managed firewalls.
  1. Before you begin, ensure that you install a License API key on the firewall.
  2. Log in to the Panorama web interface.
  3. Verify the Palo Alto Networks update server configuration for the firewalls.
     1. Select **Device** > **Setup** > **Services**.
     2. Confirm that **Update Server** is set to `updates.paloaltonetworks.com`.
     3. Confirm that **Verify Update Server Identity** is selected.
  4. Create a Deployment Profile for the new license if you have not done so already. This profile is required to generate the new authorization code for the migrated Panorama.
  5. Retrieve the VM-Series authorization code. A firewall authorization code for a flexible license begins with the letter D, as shown below.

 ©2024 Palo Alto Networks, Inc.

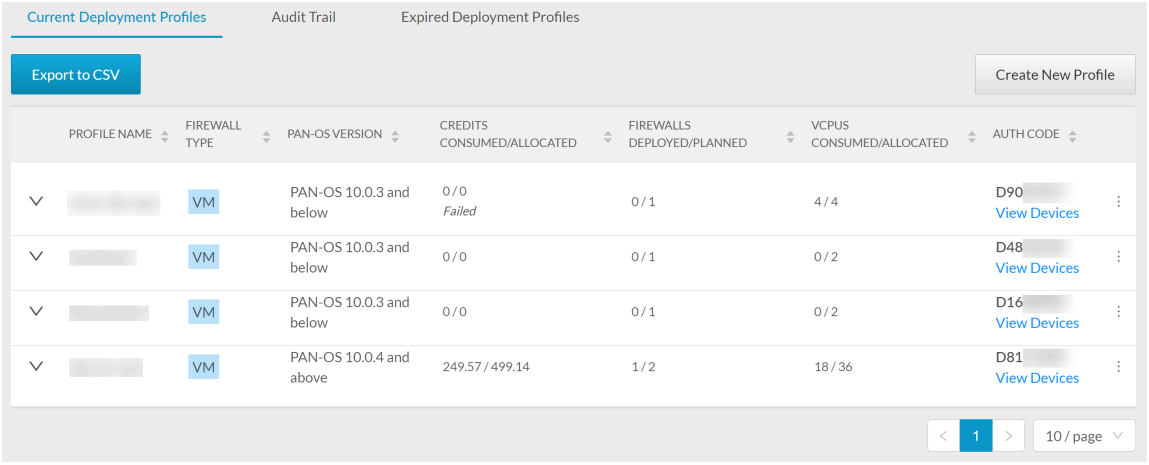

 ©2024 Palo Alto Networks, Inc.

6. Apply the new authorization code.
   1. Select **Panorama** > **Device Deployment** > **Licenses** and **Activate**.
   2. Enter your VM-Series authorization code.
   3. Use the filters to select the managed firewalls to be licensed.
   4. Enter your authorization code in the **Auth Code** column for each firewall.

 ©2024 Palo Alto Networks, Inc.

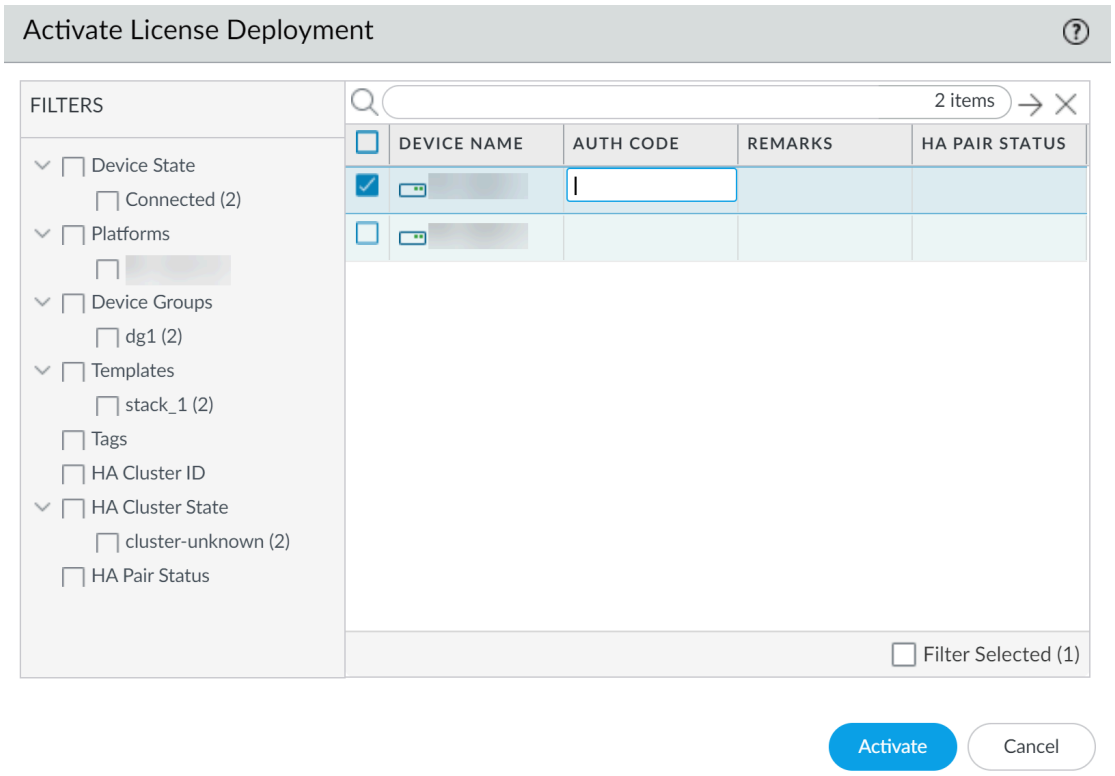

 ©2024 Palo Alto Networks, Inc.

5. **Activate** to confirm the license upgrade. Panorama contacts the Palo Alto Networks update server and consumes the tokens required for your firewalls based on the VM-Series model, vCPUs, and services you have chosen.

7. (optional) Verify the migration.

 ©2024 Palo Alto Networks, Inc.

- **Panorama cannot access the CSP**—Migrate licenses on offline Panorama managed firewalls
  1. Before you begin, ensure that you install a License API key on the firewall.
  2. Deactivate the VM from Panorama using the manual procedure, and log in to the CSP and use the token file to deactivate the VM.
  3. In the CSP, Create a Deployment Profile with the same VM-Series model, number of vCPUs, security subscriptions, and Panorama as the fixed model license. You will use the auth code from this profile.
  4. Select the new profile and click the vertical ellipsis and select **Register Firewall**.
     1. Enter the VM and firewall information and select **Submit**. This associates the firewall with the profile and its authcode and assigns a serial number.
     2. Click **View Devices** to see associated devices in **Software NGFW Devices**.
     3. In the **License** column, download the license keys to a location from which you can safely transfer the files to the host machine.
  5. Apply the new authorization code.
     1. Select **Panorama** > **Device Deployment** > **Licenses** and click **Activate**.
     2. Use the filters to select the managed firewalls to be licensed.
     3. Enter your authorization code from your deployment profile in the **Auth Code** column for each firewall.

 ©2024 Palo Alto Networks, Inc.

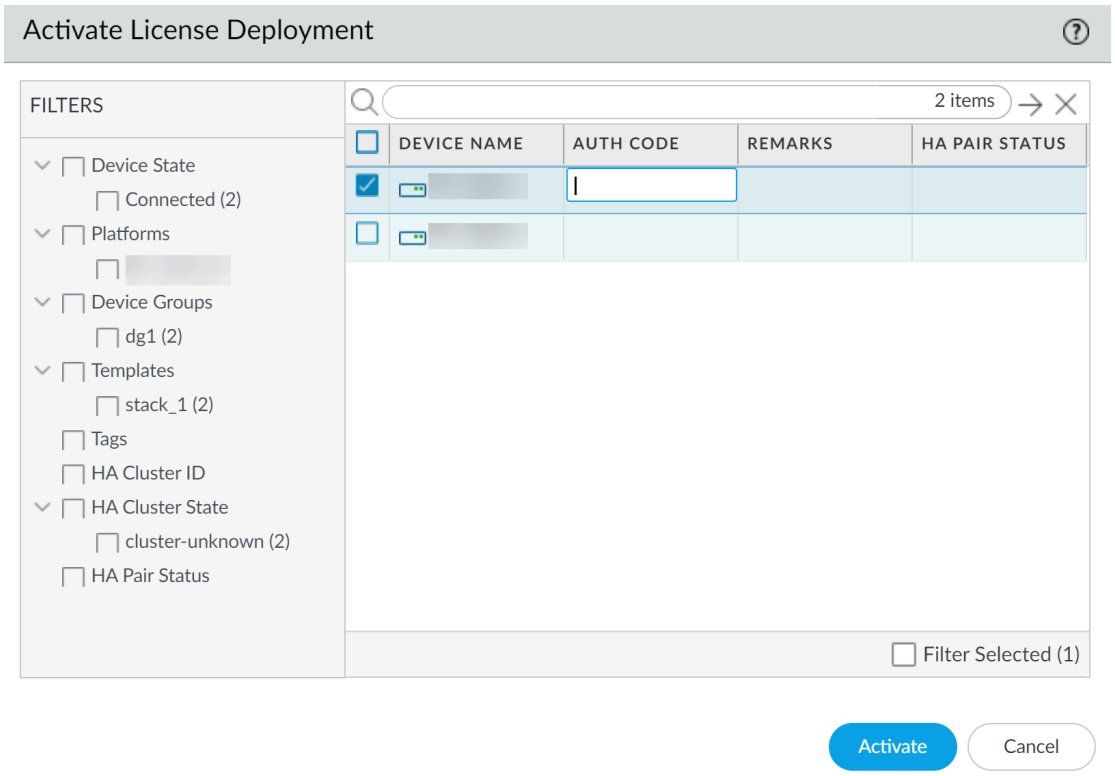

 ©2024 Palo Alto Networks, Inc.

4. Click **Activate** to confirm the license upgrade. Panorama contacts the Palo Alto Networks update server and consumes the tokens required for your firewalls based on the VM-Series model, vCPUs, and services you have chosen.

6. (optional) Verify the migration.

- Verify the migration

1. Check the license expiration date to verify the license updated successfully.
2. Verify that all subscriptions enabled in your deployment profile are applied to your device.
3. On the CSP, verify the expected number of credits allocated and credits consumed against your credit pool match.
4. On the CSP, verify that the associated tokens or quantity of licenses have been returned to your previous auth code.

 ©2024 Palo Alto Networks, Inc.

VM-Series Models

The VM-Series firewall is available in the following fixed vCPU models—VM-50, VM-100, VM-200, VM-300, VM-500, VM-700, and VM-1000-HV. These models are available for all supported PAN-OS versions, unless otherwise noted below. The software package (*.xva*, *.ova*, or *.vhdx* file) that is used to deploy the VM-Series firewall is common across all models.

You can migrate your fixed model ELA or perpetual license to a VM-Series Models *and retain the fixed model, or you can replace the license with a flexible vCPU license. See* VM-Series Models *to compare the licensing methods.*

- All models can be deployed as guest virtual machines on VMware ESXi and vCloud Air, KVM, Microsoft Hyper-V, Cisco ACI, Cisco ENCS, and Cisco CSP.
- In public cloud environments—Amazon Web Services, Azure, Google Cloud Platform, Oracle Cloud Infrastructure, Alibaba Cloud—all models except the VM-50 are supported.
- For VMware NSX, only the VM-100, VM-200, VM-300, VM-500, and VM-1000-HV firewalls are supported.

When you apply the capacity license on the VM-Series firewall, the model number and the associated capacities are implemented on the firewall. Capacity is defined in terms of the number of sessions, rules, security zones, address objects, IPSec VPN tunnels, and SSL VPN tunnels that the VM-Series firewall is optimized to handle. To make sure that you purchase the correct model for your network requirements, use the following table to understand the maximum capacity for each model and the capacity differences by model:

Model	Sessions	Security Rules	Dynamic IP Addresses	Security Zones	IPSec VPN Tunnels	SSL VPN Tunnels
VM-50	50,000	• 250 • 200 in Lite mode	1,000	15	• 250 • 25 in Lite mode	• 250 • 25 in Lite mode
VM-100 VM-200	250,000	1,500	2,500	40	1,000	500
VM-300 VM-1000-HV	800,000	10,000	100,000	40	2,000	2,000
VM-500	2,000,000	10,000	100,000	200	4,000	6,000
VM-700	10,000,000	20,000	100,000	200	8,000	12,000

©2024 Palo Alto Networks, Inc.

For information on the platforms on which you can deploy the VM-Series firewall, see VM-Series Deployments. For more information about the VM-Series firewall models, see the Palo Alto Networks Firewall comparison tool. You can also review general information About the VM-Series Firewall.

- VM-Series System Requirements
- CPU Oversubscription
- VM-50 Lite Mode
- VM-Series Model License Types
- Activate VM-Series Model Licenses
- Register the VM-Series Firewall
- Install a Device Certificate on the VM-Series Firewall
- Switch Between the BYOL and the PAYG Licenses
- Switch Between VM-Series Model Licenses
- Deactivate License(s)
- Renew VM-Series Firewall License Bundles

VM-Series System Requirements

Each instance of the VM-Series firewall requires a minimum resource allocation—number of CPUs, memory, and disk space, on its host server. Use the table below to verify that you allocate the necessary hardware resources for your VM-Series model or memory profile.

VM-Series Model	Supported Hypervisors	Supported vCPUs	Minimum Memory	Minimum Memory with GTP Enabled	Minimum Hard Drive
VM-50	ESXi, Hyper-V, KVM	2	• 5.5GB • 4.5GB in Lite mode	• 6GB • 5GB in Lite mode	32GB (60GB at boot)
VM-100	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, Cisco ENCS, NSX-T (VM-100)	2	6.5GB	7.5GB	60GB

 ©2024 Palo Alto Networks, Inc.

VM-Series Model	Supported Hypervisors	Supported vCPUs	Minimum Memory	Minimum Memory with GTP Enabled	Minimum Hard Drive
	The VM-100 on Azure requires 4 vCPUs.				
VM-300	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, Cisco ENCS, NSX-T (VM-300)	2, 4	9GB*	10GB	60GB
VM-500	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, NSX-T	2, 4, 8	16GB	20GB	60GB
VM-700	AWS, Azure, ESXi, Google Cloud Platform, Hyper-V, KVM, OCI, Alibaba Cloud, Cisco ACI, Cisco CSP, NSX-T	2, 4, 8, 16	56GB	64GB	60GB

*In 10.0.X, 9 GB might be insufficient depending upon the feature set or combination of feature sets (such as GTP or high-performance features) used on the firewall. If you experience memory resource related issues, increase memory to 11 GB to accommodate the additional memory requirements of some of the features or combination of features. Alternately, you can Enable ZRAM on the VM-Series Firewall to improve memory usage.

You can enable Lite mode on the VM-50. Lite mode is an alternative operating mode for environments where resources are limited. See VM-50 Lite Mode for more information.

To achieve the best performance, all of the needed cores should be available on a single CPU socket.

For operation, the VM-50 firewall requires minimum 32GB of hard drive space. However, because the VM-Series base image is common to all models, you must allocate 60GB of hard drive space until you license the VM-50.

The number of vCPUs assigned to the management plane and those assigned to the dataplane differs depending on the total number of vCPUs assigned to the VM-Series firewall. If you assign

 ©2024 Palo Alto Networks, Inc.

more vCPUs than those officially supported by the license, any additional vCPUs are assigned to the management plane.

Total vCPUs	Management Plane vCPUs	Dataplane vCPUs
2	1	1
4	2	2
8	2	6
16	4	12

CPU Oversubscription

The VM-Series firewall supports CPU oversubscription on all models. CPU oversubscription allows you deploy a higher density of VM-Series firewalls on hypervisors running on x86 architecture. You can deploy two (2:1) to five (5:1) VM-Series firewalls per required allocation of CPUs. When planning your deployment, use the following formula to calculate the number of VM-Series firewalls your hardware can support.

(Total CPUs x Oversub Ratio)/CPUs per firewall = total number of VM-Series firewalls

For example, at a 5:1 ratio, a host machine with 16 physical CPU and at least 180GB of memory (40 × 4.5GB) can support up to 40 instances to the VM-50. Each VM-50 requires two vCPUs and five VM-50s can be associated to each pair of vCPUs.

(16 CPUs x 5)/2 = 40 VM-50 firewalls

Beyond meeting the minimum VM-Series System Requirements, no additional configuration is required to take advantage of oversubscription. Deploy VM-Series firewalls normally and resource oversubscription occurs automatically. When planning your deployment, consider other functions, such as virtual switches, and guest machines on the host that require hardware resources of their own.

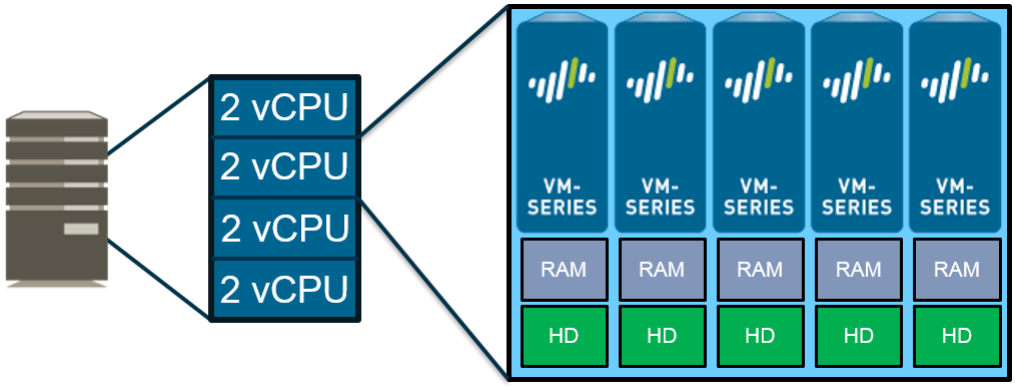

 ©2024 Palo Alto Networks, Inc.

VM-50 Lite Mode

The standard VM-50, while the smallest model of the VM-Series, requires more resources than are available in some environments. The VM-50 Lite mode provides an alternative for environments where hardware resources are constrained. The VM-50 Lite requires 4.5GB of memory instead of the 5.5GB required by the standard VM-50. The VM-50 Lite uses the same license as the standard VM-50 but comes up in Lite mode when allocated 4.5GB of RAM.

- *In high availability deployments, both VM-Series firewalls must both be licensed as a VM-50 Lite to avoid capacity mismatch issues. In the case of a capacity license mismatch, the VM-50 (non-Lite) is considered to have a higher capacity; the VM-50 becomes non-functional while the VM-50 Lite remains functional.*
- *The VM-50 Lite does not support jumbo frames; the VM-50 and VM-50 Lite do not support WildFire inline ML.*

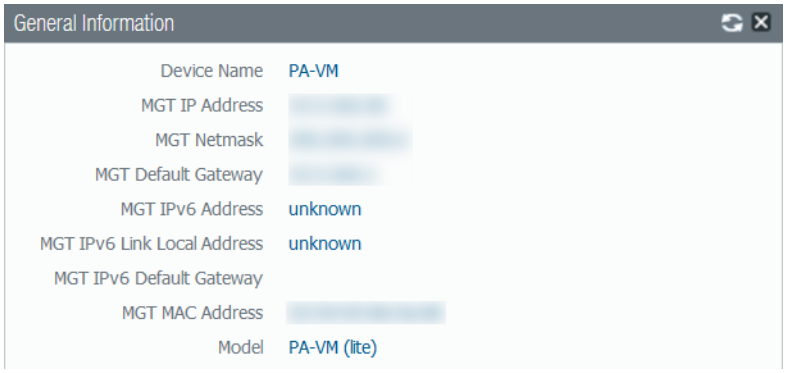

VM-Series Model License Types

New capacity licenses (non-Software NGFW Credits) are no longer available for purchase. However, you one (1) year renewals for capacity (perpetual and term-based) licenses are available.

The following licenses and subscriptions are available for the VM-Series firewall:

- **Capacity License**—The VM-Series firewall requires a base license, also called a *capacity license*, to enable the model number (VM-50, VM-100, VM-200, VM300, VM-500, VM-700, or VM-1000-HV) and the associated capacities on the firewall. Capacity licenses are included in a bundle and can be perpetual or term-based:
  - **Perpetual License**—A license with no expiration date, it allows you to use the VM-Series firewall at the licensed capacity, indefinitely. Perpetual licenses are available for the VM-Series capacity license only.
  - **Term-Based License**—A term-based license allows you to use the VM-Series firewall for a specified period of time. It has an expiration date and you will be prompted to renew the license before it expires. Term-based licenses are available for the capacity licenses, support entitlements, and subscriptions.

 ©2024 Palo Alto Networks, Inc.

- **VM-Series ELA**—For high-growth enterprises, the VM-Series enterprise licensing agreement (VM-Series ELA) provides a fixed price licensing option that allows up to unlimited deployment of VM-Series firewalls with BYOL. The ELA is offered in one and three-year term agreements with no true-up at the end of the term.

  There are two flavors of the VM-Series ELA:

  - If you purchased the VM-Series ELA before December 4, 2018, you have the legacy VM-Series ELA which includes your choice of a single VM-Series model on any supported hypervisor or public cloud environment. With this ELA, you receive a single license authorization code for capacity, support, GlobalProtect, PAN-DB URL Filtering, Threat Prevention, WildFire subscriptions for every instance of the VM-Series firewall. You also get unlimited deployments of the Panorama virtual appliance included with a device management license for 1000 firewalls on each.

    Palo Alto Networks began phasing out the legacy VM-Series ELA on April 16, 2019. Existing enterprise license customers will be notified by their support representative when their account is migrated to the Multi-Model ELA. Licensing tokens will be distributed according to your VM-Series firewall subscription agreement — no additional action is necessary for continued operation of your firewalls. If you would like to Manage VM-Series ELA License Tokens, you must designate an ELA administrator. Only a super user role on the Palo Alto Networks Customer Support Portal (CSP) can assign an ELA administrator.
  - The VM-Series Enterprise License Agreement (Multi-Model ELA) you purchase after December 4, 2018 (either as a new purchase or as a repurchase of the legacy VM-Series ELA) is called the multi-model VM-Series ELA that includes most models of the VM-Series firewall portfolio along with the GlobalProtect, PAN-DB URL Filtering, Threat Prevention, WildFire subscriptions, and support entitlement. You also get unlimited deployments of the Panorama virtual appliance with a device management license for 1000 firewalls on each.

VM-Series Firewall Licenses for Public Clouds

The VM-Series firewall licensing strategy is the same for AWS, Azure, and Google Cloud Platform. There are different license types (see License Types—VM-Series Firewalls), and Bring Your Own License and Pay-as-you-go licensing methods:

- **Bring Your Own License (BYOL)**—A license that is purchased from a partner, reseller, or directly from Palo Alto Networks. BYOL supports individual capacity licenses, support licenses, and subscription bundles.
  - For individual BYOL licenses, you must apply the auth code after you deploy the VM-Series firewall.
  - A BYOL license bundle has a single auth code you can include in the bootstrap package (see Bootstrap the VM-Series Firewall). All the subscriptions included in the bundle are licensed when the firewall launches.

A BYOL license for the VM-Series firewall on OCI GovCloud requires PAN-OS 10.0.5 or later for FIPS and non-FIPS modes.

 ©2024 Palo Alto Networks, Inc.

- **Pay-as-you-go (PAYG)**—Also called *usage-based* or *pay-per-use* licensing. PAYG licenses can be purchased from your Cloud provider:
  - AWS: Purchase from AWS Marketplace. Supports hourly and annual PAYG options.
  - Azure: Purchase from Azure Marketplace. Supports the hourly PAYG option.
  - Google Cloud Platform: Purchase from Google Cloud Platform Marketplace. Google Cloud Platform supports per-minute PAYG option.
  - Oracle Cloud Infrastructure: (PAN-OS 10.0.3 or later) Purchase from Oracle Cloud Marketplace.

    The VM-Series on OCI PAYG license does not support the VM-100.

  With the PAYG license bundles, the firewall is prelicensed and ready for use as soon as you deploy it; you do not receive an auth code. When you stop or terminate the firewall from your Cloud console, PAYG licenses are suspended or terminated.

  A PAYG license applies a VM-Series capacity license based on the hardware allocated to the instance. the PAYG instance checks the amount of hardware resources available to the instance and applies the largest VM-Series firewall capacity license allowed for the resources available. For example, if the instance has 2 vCPUs and 16GB of memory, a VM-100 capacity license is applied based on the number of vCPUs. However, if the instance has 16 vCPUs and 16GB of memory, a VM-500 license is applied based on the amount of memory. For more information about VM-Series model resource requirements, see VM-Series System Requirements.

  Downgrading PAN-OS is not supported on a PAYG firewall instance that was initially deployed running PAN-OS 9.1.2. Firewall instances deployed prior to PAN-OS 9.1.2 can be downgraded to older versions of PAN-OS.

  The PAYG licenses are bundled as follows:

License Features	Bundle 1	Bundle 2
VM-Series firewall capacity license	VM-100, VM-300, VM-500, VM-700	VM-100, VM-300, VM-500, VM-700
Premium Support	✓	✓
Threat Prevention (AV, IPS, and malware prevention)	✓	✓
GlobalProtect		✓
PAN-DB URL Filtering		✓
WildFire		✓
DNS Security		✓

 ©2024 Palo Alto Networks, Inc.

When using the VM-Series firewall CLI to view your applied PAYG license, the command show system info displays a different value from the output displayed for the command ***request license info****. The command* ***request license info*** *always displays the license as a VM-300, regardless of the actual VM-Series model.*

You cannot switch between the PAYG and the BYOL licenses. To move from PAYG to BYOL, contact your Palo Alto Networks channel partner or sales representative to purchase a BYOL license and get a BYOL auth code that you can use to license your firewall. If you have deployed your firewall and want to switch the license, see Switch Between the BYOL and the PAYG Licenses.

If you have an evaluation copy of the VM-Series firewall and would like to convert it to a fully licensed (purchased) copy for the same license type (BYOL to BYOL), you can deactivate the evaluation license and activate the purchased license in its place. See Upgrade the VM-Series Firewall for instructions.

VM-Series Enterprise License Agreement (Multi-Model ELA)

The VM-Series Enterprise License Agreement (VM-Series ELA) is a one- or three-year comprehensive licensing agreement that enables you to purchase VM-Series firewalls, along with the GlobalProtect, PAN-DB URL Filtering, Threat Prevention, WildFire, and DNS Security subscriptions. It also includes a support entitlement and a device management license for Panorama. The multi-model VM-Series ELA provides simplified license management with a single contract that allows you to deploy any model of the VM-Series firewall that meets your enterprise security needs.

When you purchase the multi-model VM-Series ELA, you forecast the number of firewalls that you'll need over the term of your subscription. Based on your forecast and an additional allotment that accommodates for future growth, your account on the Customer Support Portal (CSP) is credited with a license token pool that allows you to deploy any model of the VM-Series firewall. Depending on the firewall model and the number of firewalls that you deploy, a specified number of tokens are deducted from your available license token pool. The tokens drawn from your account are calculated based on the value of each firewall model:

- VM-50—10 tokens
- VM-100—25 tokens
- VM-300—50 tokens
- VM-500—140 tokens
- VM-700—300 tokens

With the VM-Series ELA, there is no true-up due at the end of the term which means that you are not billed retroactively even if you deploy more firewalls than your original forecast. So, to balance flexibility with accountability, the VM-Series ELA terms of use include a bounded and unbounded period that explains how you can consume tokens and deploy firewalls as the need arises. For details, refer to the ELA terms and conditions. The VM-Series firewalls that you deploy with the VM-Series ELA do not have a perpetual license and on the expiry of the term, you must renew the agreement to extend the support entitlement and get continued access to software and content release updates on the firewalls.

With the ELA administrator role on the CSP, you can transfer or split the licensing tokens among other administrators who belong to different departments with their own CSP accounts. This

 ©2024 Palo Alto Networks, Inc.

sharing enables other administrators in your enterprise to deploy the VM-Series firewall on demand as long as they have tokens available in their respective CSP accounts. See Manage VM-Series ELA License Tokens to invite other administrators to share ELA tokens and deploy any model of the VM-Series firewall that meets your enterprise security need. You can also reclaim tokens to remove CSP accounts from the VM-Series ELA if you want to redistribute tokens based on changing organizational needs.

Watch VM-Series Multi-Model ELA videos

- Manage VM-Series ELA License Tokens
- Accept the VM-Series ELA

Manage VM-Series ELA License Tokens

The VM-Series Enterprise License Agreement (Multi-Model ELA) (VM-Series ELA) gives you the flexibility of having a single contract that you can share with other administrators in your enterprise. You must have the super user role on the Palo Alto Networks Customer Support Portal (CSP) to activate the ELA, and upon activating the ELA authorization code you inherit the ELA administrator role on the CSP.

With the ELA administrator role, you can manage the license token pool available to deploy VM-Series firewalls and subscriptions included in the agreement. You can invite other administrators to share the VM-Series ELA tokens, grant which models and how many instances of the VM-Series firewalls are available to each administrator, as well as remove CSP accounts from your VM-Series ELA. Depending on what you allocate for each grantee, they receive a specific number of tokens that they can then use to deploy VM-Series firewalls.

Additional purchases and grants do not directly add to the number of available VM-Series firewalls in a CSP account; instead, ELA license tokens are added to the VM-Series ELA token pool. The ELA license tokens can subsequently be allocated by the ELA administrator to a given CSP account to increase the number of available VM-Series firewalls.

STEP 1 | (Legacy VM-Series ELA Customers only) Designate an ELA administrator to manage tokens.

Existing enterprise license customers who have been migrated to the Multi-Model ELA must designate an ELA administrator to manage VM-Series ELA license tokens. Upon conversion, no other action is necessary for continued operation of your firewalls, however, you will not be able to (re)allocate tokens for deploying firewalls until an ELA administrator has been assigned. Only an administrator with a super user role on the CSP has the ability to designate an ELA administrator, who in turn, can manage tokens or grant tokens to other administrators.

1. Log in to the Palo Alto Networks CSP.
2. Select **Members** > **Manage Users**.
3. Click on the pencil icon under **Actions** to edit the user to whom you want to assign the ELA administrator role.
4. Select **ELA Administrator** and then click the check mark to add the new role to the selected user.
5. Continue to step 3.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Activate the ELA authorization code.

The administrative user who activates the ELA inherits the ELA administrator and super user role on the CSP and has the ability to manage the tokens or grant the tokens to other administrators.

1. Log in to the Palo Alto Networks CSP.
2. Select **Assets** > **Enterprise Agreements** > **Activate Enterprise Agreement**.
3. Enter the **Authorization Code** and **Agree and Submit** the EULA.

   Verify the authorization code is registered to your account under Enterprise Agreements: VM-Series. The page displays the Auth Code, Account ID, Account Name, License Description, Expiration Date, the number of Licenses (used/total) you have, and

 ©2024 Palo Alto Networks, Inc.

how many are available to deploy within the bounded and unbounded period of the agreement.

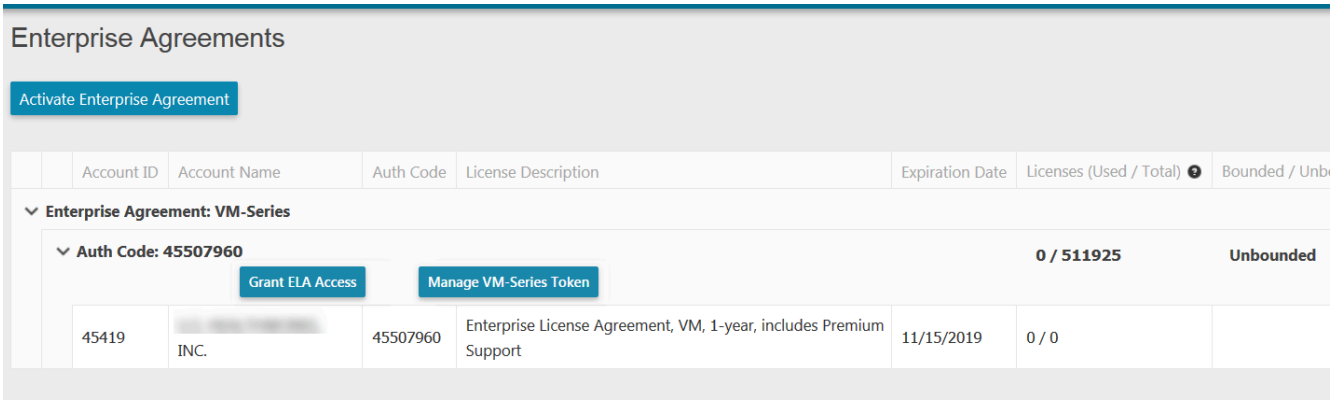

4. Select **Assets** > **VM-Series Auth-Codes** to view the authorization codes for deploying each model of the VM-Series firewall and associated subscriptions included with the ELA.

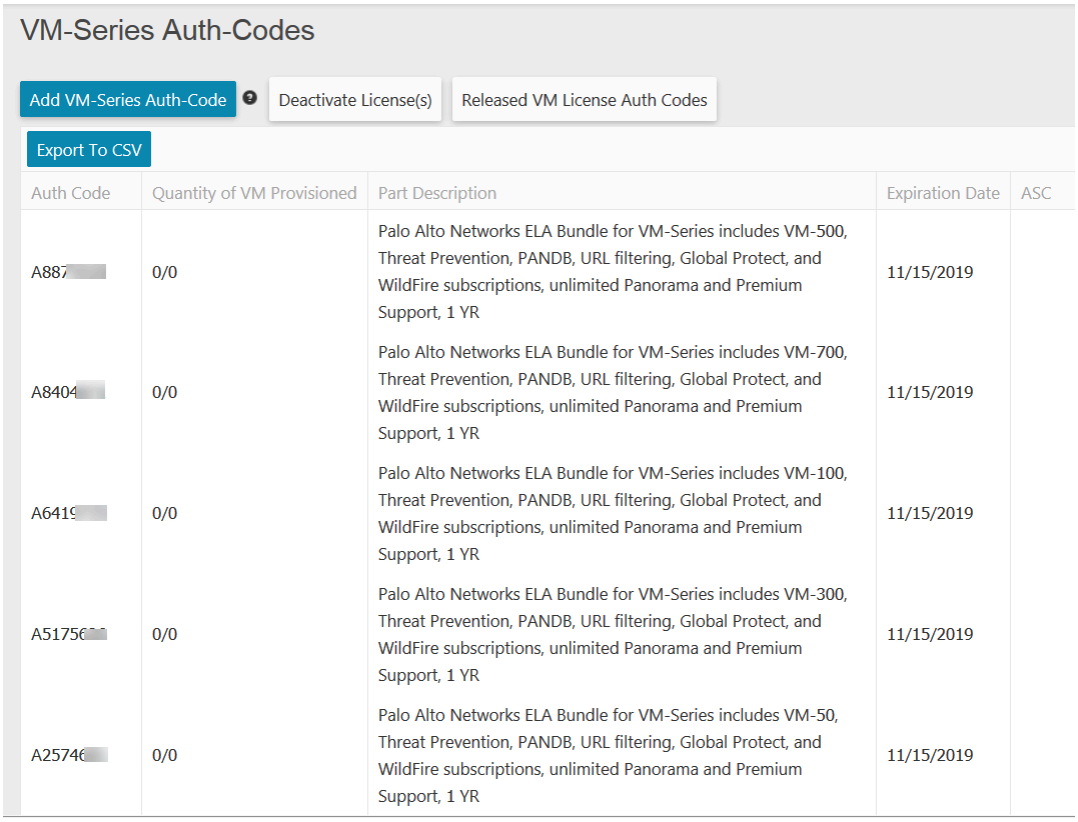

STEP 3 | Grant ELA access to other administrators in your enterprise.

This capability allows you to share the VM-Series ELA with other administrators within your enterprise or department so that they can deploy VM-Series firewalls on demand. As an ELA

 ©2024 Palo Alto Networks, Inc.

administrator, you can grant access to other users who are registered with an email address on the CSP.

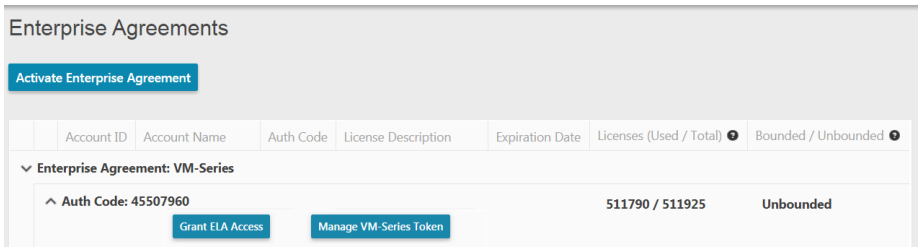

1. On **Assets** > **Enterprise Agreements**, select **Grant ELA Access**.
2. Enter the **Destination Email** address of the administrator whom you want to invite.

   The destination email address that you enter above must be a registered user on the CSP with a super user role so that they can log in and accept the grant. If the email address is not registered on the CSP, you must first create a new account for the user on **Members** > **Create New User**.
3. Select **Notify User** to trigger a notification email to the email address you entered.

   The recipient must log in to the CSP to Accept the VM-Series ELA. After the recipient accepts the grant, the account ID is available on **Assets** > **Enterprise Agreements** as shown in the following screenshot.

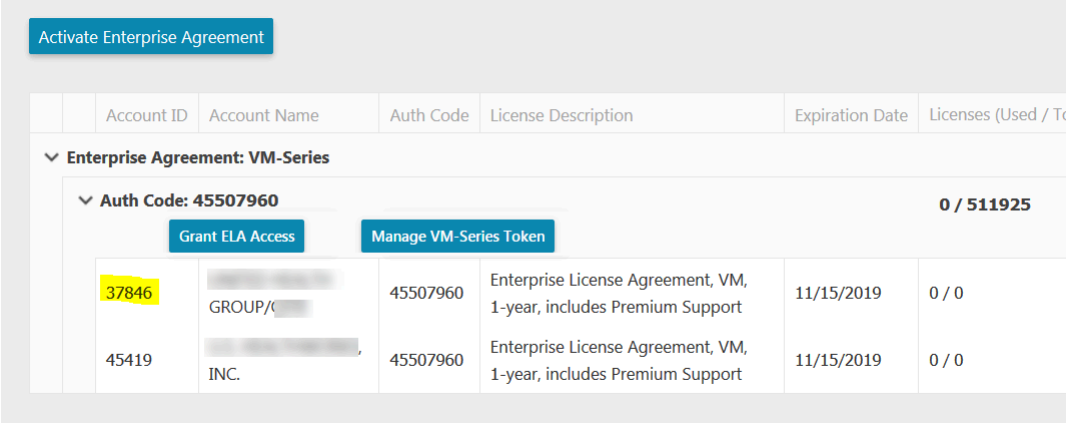

STEP 4 | Allocate tokens for deploying firewalls.

1. Select **Assets** > **Enterprise Agreements** > **Manage VM-Series Tokens**.

   For each account ID, you can specify the number of firewalls by model that you want to allocate. Based on the quantity and firewall model, the number of tokens are

 ©2024 Palo Alto Networks, Inc.

automatically calculated and become available for use. In this example, you are allowing 10 instances each of the VM-50 and the VM-500.

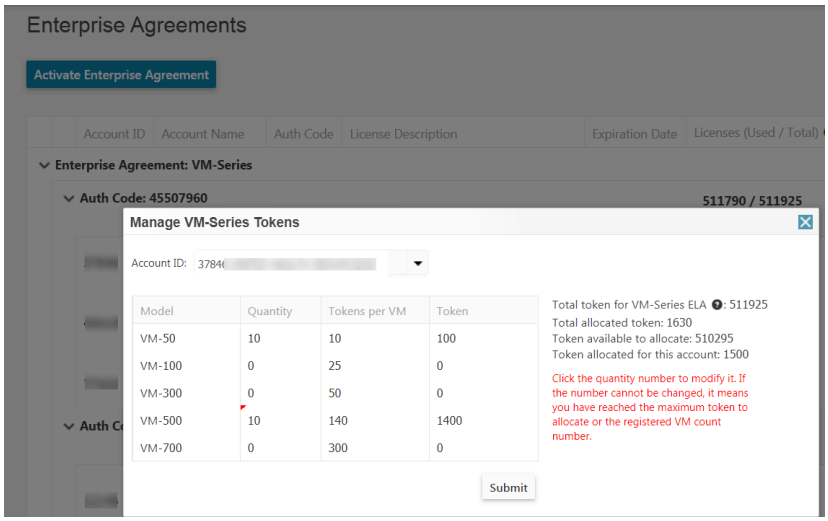

2. Verify that the accurate number of firewall instances are deposited in the account.

   Select **Assets** > **VM-Series Auth-Code** to confirm the auth codes you allotted. In this example, the account has the ability to provision 10 instances each of the VM-50 and the VM-500. As the recipients deploy firewalls, the number of tokens are deducted from the total available pool, and you can view the number of firewall instances that they have provisioned as a ratio of the total quantity you allocated for them. As your security needs evolve, you have the flexibility to allocate more quantity and allow access to a different VM-Series firewall model as long as you have tokens available.

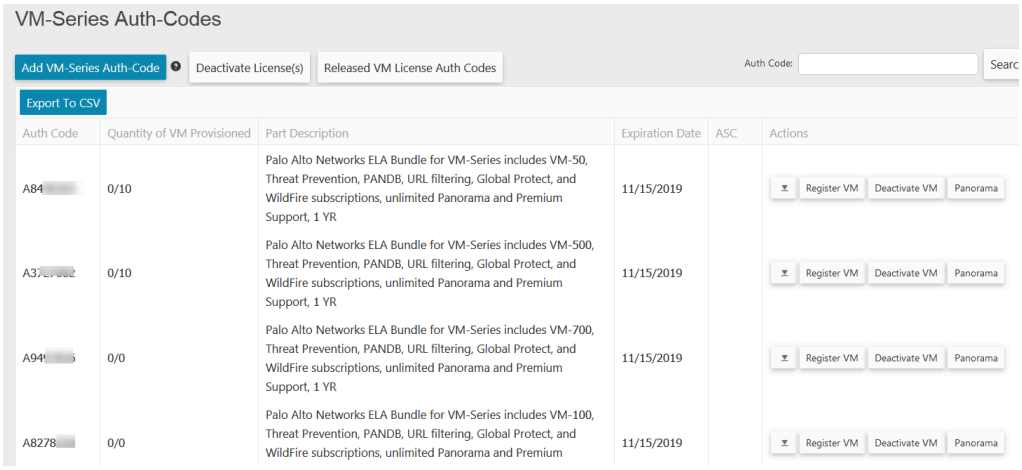

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Remove a CSP account from the VM-Series ELA to reclaim tokens.

You cannot reclaim a portion of the tokens allocated to a CSP account. By reclaiming tokens, you are removing the entirety of the CSP account from the VM-Series ELA and reallocating all associated tokens to the token pool.

1. Verify that all tokens associated with the CSP account that you want to remove are not being utilized by the VM-Series firewalls. Deactivate the VM-Series firewalls as necessary to provision tokens for removal.
2. Select **Assets** > **Enterprise Agreements** > **Manage VM-Series Token**.

   Select the account ID from whom you want to reclaim tokens from and click **Reclaim Token**. If tokens are available for reclamation, you will receive a confirmation of a successful removal.

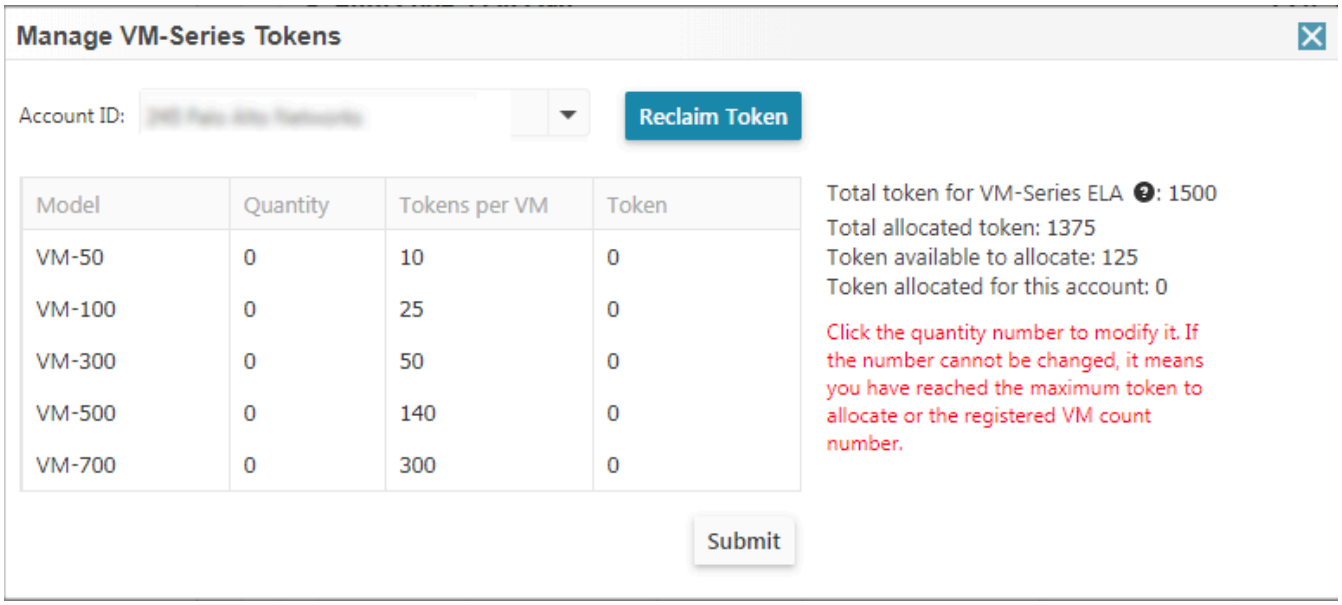

Accept the VM-Series ELA

If your enterprise has purchased a VM-Series ELA, your ELA administrator can invite you to share the contract and share the license token pool so that you have access to VM-Series firewall auth codes which enable you to deploy VM-Series firewalls on demand. When you receive a grant for access to the VM-Series ELA, you get an email notification that includes a link to log in to the Palo Alto Networks Customer Support Portal (CSP) and you must agree and accept the terms of use. After you accept the ELA terms of use, the ELA administrator can allocate which VM-Series firewall models and how many you are entitled to use; the corresponding number of VM-Series ELA tokens are deposited in your account.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Check your email inbox for the grant notification.

The notification includes the email address of the ELA administrator who has invited you to share the VM-Series ELA.

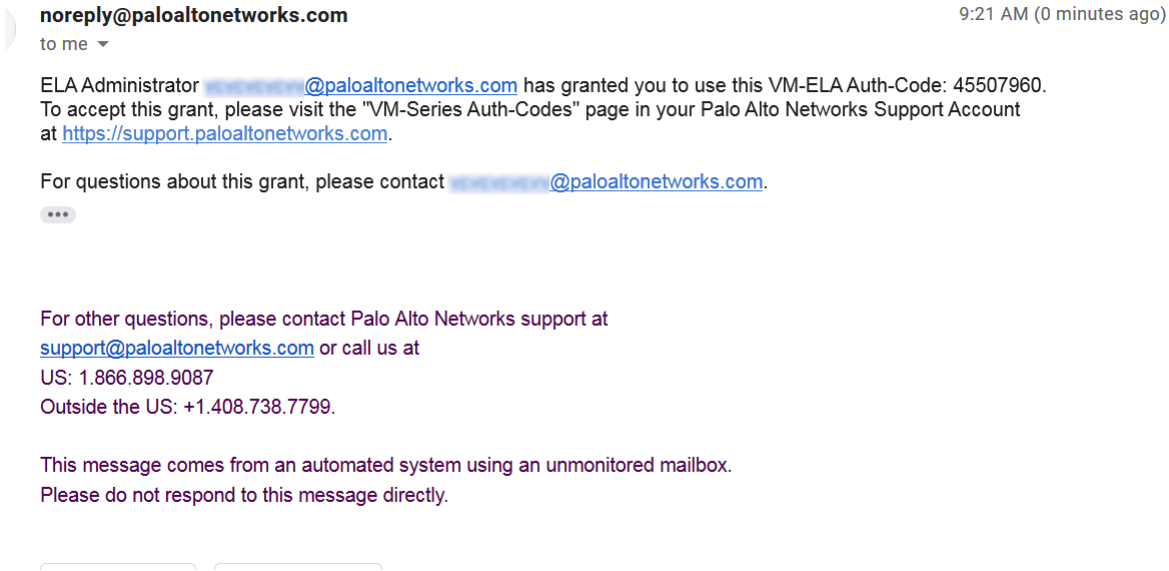

STEP 2 | Accept the grant.

You must review the terms and accept the EULA and the support agreement before the ELA admin can allocate tokens which enable you to deploy VM-Series firewalls.

1. Log in to the Palo Alto CSP.
2. Select **VM-Series Auth Codes** to **Review Tokens Grant**.

   You must Agree and Accept the EULA and support agreement to accept the grant. If you reject it, the ELA Admin who gave you the grant receives an email notification that you declined the grant. Do make sure to let the ELA administrator know that you have

 ©2024 Palo Alto Networks, Inc.

accepted the grant so that you he/she can allocate the VM-Series firewall models and quantity that you can deploy.

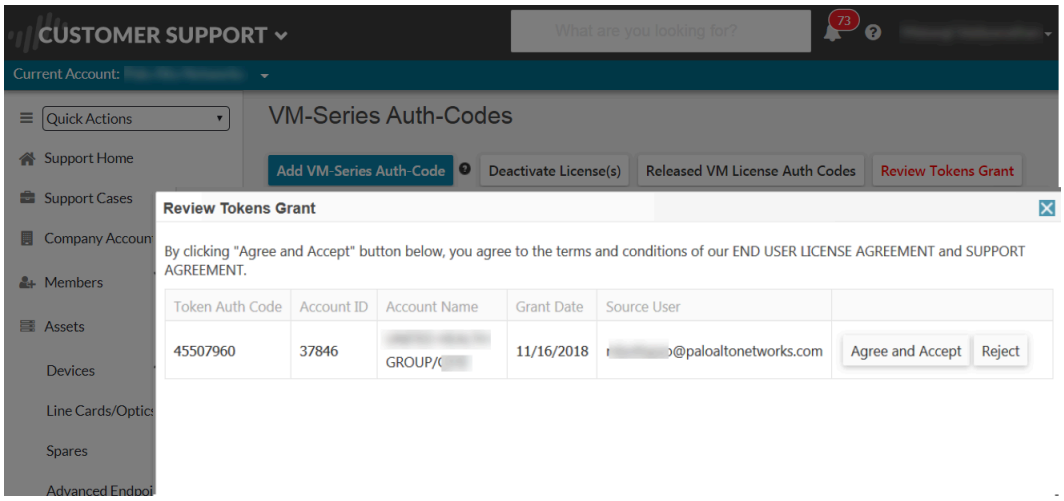

If you belong to multiple accounts on the CSP and accidentally accept the grant in to the wrong account, you must request the ELA administrator to resend the grant to you. Do not start using the auth code to provision firewalls until you accept the grant in the correct account.

STEP 3 | Verify which VM-Series models and how many are allocated for you.

After the ELA administrator allocates the VM-Series firewall models and number of instances you can provision, you can select **Assets** > **VM-Series Auth Codes** to view which models and how many of each are allocated for you. For example, the grant in the following screenshot displays the auth codes that enable you to deploy 10 instances each of the VM-50 and the VM-500.

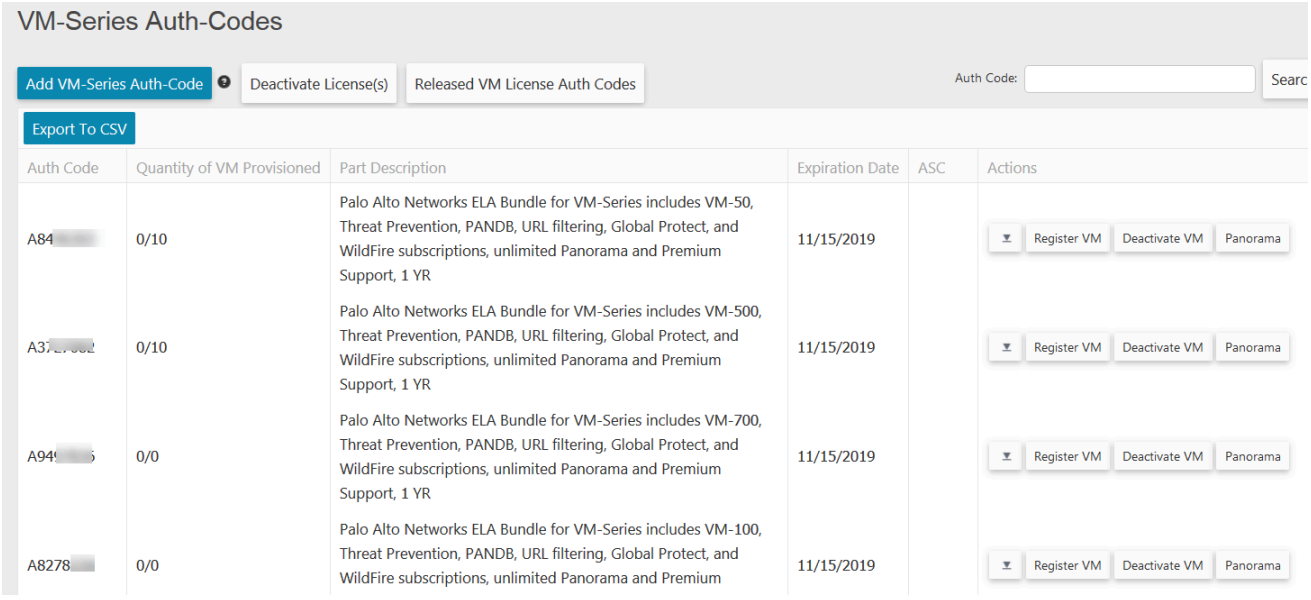

As you deploy firewalls and register them to the CSP, the number of provisioned firewalls is incremented. The **Quantity of VM Provisioned** displays the ratio of provisioned to total available for each model.

 ©2024 Palo Alto Networks, Inc.

Activate VM-Series Model Licenses

To activate the license on your VM-Series firewall, you must have deployed the VM-Series firewall and completed initial configuration. To deploy the firewall, see VM-Series Deployments.

Use the instructions in this section for all the BYOL models including AWS, Azure, and Google Public Cloud. For usage-based licensing in public clouds, you do not need to activate the license. You must Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code) in order to activate your premium support entitlement.

For usage-based models of the VM-Series firewall in the AWS Marketplace, instances with short and long AWS instance IDs are supported.

Until you activate the license on the VM-Series firewall, the firewall does not have a serial number, the MAC address of the dataplane interfaces are not unique, and only a minimal number of sessions are supported. Because the MAC addresses are not unique until the firewall is licensed, to prevent issues caused by overlapping MAC addresses, make sure that you do not have multiple, unlicensed VM-Series firewalls.

When you activate the license, the licensing server uses the UUID and the CPU ID of the virtual machine to generate a unique serial number for the VM-Series firewall. The capacity auth code in conjunction with the serial number is used to validate your entitlement.

*The VM-Series firewall **License** tab displays a standard VM-300 license file for all license models. To find your specific license model information, view the system info in the UI, or the use the CLI to view **`system info`**.*

After you license a VM-Series firewall, if you need to delete and redeploy the VM-Series firewall, make sure to Deactivate the License(s) on the firewall. Deactivating the license allows you to transfer the active licenses to a new instance of the VM-Series firewall without help from technical support.

- Activate the License for the VM-Series Firewall (Standalone Version)
- Activate the License for the VM-Series Firewall for VMware NSX
- Troubleshoot License Activation Issues

Activate the License for the VM-Series Firewall (Standalone Version)

If you have not elected to use the bootstrapping workflow using a subscription bundle, you must deploy the VM-Series firewall and complete initial configuration before you can activate the license on your VM-Series firewall.

- Direct internet access
- No internet access

 ©2024 Palo Alto Networks, Inc.

- Direct internet access

  To activate the license, the firewall must be configured with an IP address, netmask, default gateway, and DNS server IP address.

  The firewall must have a valid DNS configuration and have network connectivity to access the Palo Alto Networks licensing server.

  1. Select **Device** > **Licenses** and select the **Retrieve license keys from license server** link.
  2. The firewall will connect to the update server (updates.paloaltonetworks.com), and download the license and reboot automatically.
  3. Log back in to the web interface and confirm that the **Dashboard** displays a valid serial number. If the term **Unknown** displays, it means the device is not licensed.
  4. On **Device** > **Licenses**, verify that **PA-VM** license is added to the device.

     If you see an error message, check Troubleshoot License Activation Issues.

- No internet access
  1. Select **Device** > **Licenses** and click the **Activate Feature using Auth Code** link.
  2. Click **Download Authorization File**, and download the **authorizationfile.txt** on the client machine.
  3. Copy the **authorizationfile.txt** to a computer that has access to the internet and log in to the support portal. Click **My VM-Series Auth-Codes** link and select the applicable auth code from the list and click the **Register VM** link.
  4. On the **Register Virtual Machine** tab upload the authorization file. Select the PAN-OS version and the hypervisor on which you have deployed the firewall, to complete the

 ©2024 Palo Alto Networks, Inc.

registration process. The serial number of your VM-Series firewall will be attached to your account records.

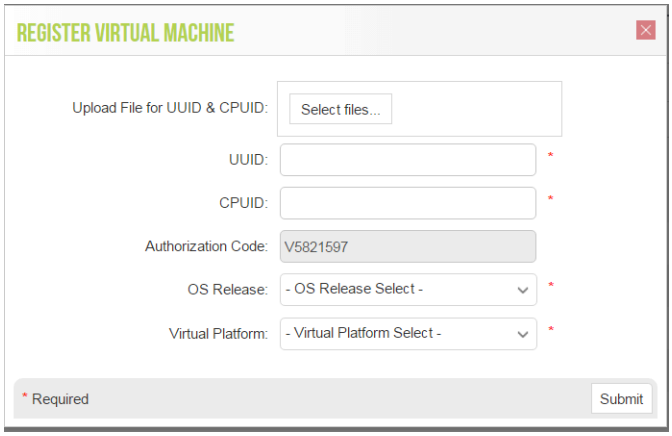

5. Navigate to **Assets** > **My Devices** and search for the VM-Series device just registered and click the **PA-VM** link. This will download the VM-Series license key to the client machine.
6. Copy the license key to the machine that can access the web interface of the VM-Series firewall and navigate to **Device** > **Licenses**.

   License keys must be installed through the web interface. The firewall does not support license key installation through SCP or FTP.
7. Click **Manually Upload License** link and enter the license key. When the capacity license is activated on the firewall, a reboot occurs.
8. Log in to the device and confirm that the **Dashboard** displays a valid serial number and that the **PA-VM** license displays in the **Device** > **Licenses** tab.

Activate the License for the VM-Series Firewall for VMware NSX

Panorama serves as the central point of administration for the VM-Series firewalls for VMware NSX and the license activation process is automated when Panorama has direct internet access. Panorama connects to the Palo Alto Networks update server to retrieve the licenses, and when a new VM-Series firewall for NSX is deployed, it communicates with Panorama to obtain the license. If Panorama is not connected to the internet, you need to manually license each instance of the VM-Series firewall so that the firewall can connect to Panorama. For an overview of the components and requirements for deploying the VM-Series firewall for NSX, see VM-Series for NSX Firewall Overview.

For this integrated solution, the auth code (for example, PAN-VM-1000-HV-SUB-BND-NSX2) includes licenses for threat prevention, URL filtering and WildFire subscriptions and premium support for the requested period.

In order to activate the license, you must have completed the following tasks:

- Registered the auth code to the support account. If you don't register the auth code, the licensing server will fail to create a license.

 ©2024 Palo Alto Networks, Inc.

- Entered the auth code in the Service Definition on Panorama. On Panorama, select **VMware Service Manager** to add the **Authorization Code** to the **VMware Service Definition**.

  If you have purchased an evaluation auth code, you can license up to 5 VM-Series firewalls with the VM-1000-HV capacity license for a period of 30 or 60 days. Because this solution allows you to deploy one VM-Series firewall per ESXi host, the ESXi cluster can include a maximum of 5 ESXi hosts when using an evaluation license.

The following process of activating the licenses is manual. If you have a custom script or an orchestration service, you can use the Licensing API to automate the process of retrieving the licenses for the VM-Series firewalls.

- Activate Licenses on VM-Series Firewalls on NSX When Panorama has Internet Access
- Activate Licenses on VM-Series Firewalls on NSX When Panorama has No Internet Access

Activate Licenses on VM-Series Firewalls on NSX When Panorama has Internet Access

Complete the following procedure to activate the VM-Series firewall for NSX when Panorama has access to the internet.

STEP 1 | Verify that the VM-Series firewall is connected to Panorama.

1. Log in to Panorama.
2. Select **Panorama** > **Managed Devices** and check that the firewall displays as Connected.

STEP 2 | Verify that each firewall is licensed.

Select **Panorama** > **Device Deployment** > **Licenses** and verify that Panorama has matched the auth code and applied the licenses to each firewall.

If you do not see the licenses, click **Refresh**. Select the VM-Series firewalls for which to retrieve subscription licenses and click **OK**.

Activate Licenses on VM-Series Firewalls on NSX When Panorama has No Internet Access

Complete the following procedure to activate the VM-Series firewall for NSX when Panorama does not have access to the internet.

STEP 1 | Locate the CPU ID and UUID of the VM-Series firewall.

1. From the vCenter server obtain the IP address of the firewall.
2. Log into the web interface and select **Dashboard**.
3. Get the **CPU ID** and the **UUID** for the firewall from the General Information widget.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Activate the auth code and generate the license keys.

1. Log in to the Palo Alto Networks Customer Support website with your account credentials. If you need a new account, see Create a Support Account.
2. Select **Assets** > **VM-Series Auth Codes**, click **Add VM-Series Auth Codes** to enter the auth code.
3. Select **Register VM** in the row that corresponds to the auth code that you just registered, enter the CPU ID and the UUID of the firewall and click **Submit**. The portal will generate a serial number for the firewall.
4. Select **Assets** > **Devices** and search for the serial number.
5. Click the link the Actions column to download each key locally to your laptop. In addition to the subscription license key, you must get the capacity license and the support license keys.

STEP 3 | Upload the keys to the firewall.

1. Log in to the firewall web interface.
2. Select **Device** > **Licenses**, and select **Manually upload license key**.
3. **Browse** to select a key and click **OK** to install the license on the firewall.

   Install the capacity license key file (pa-vm.key) first. When you apply the capacity license key, the VM-Series firewall will reboot. On reboot, the firewall will have a serial number that you can use to register the firewall as a managed device on Panorama.
4. Repeat the process to install each key on the firewall.
5. Select **Dashboard** and verify that you can see the **Serial #** in the General Information widget.

STEP 4 | Add the serial number of the firewall on Panorama.

Select **Panorama** > **Managed Devices** and click **Add** to enter the serial number for the VM-Series firewall for NSX. The firewall should now be able to connect with Panorama so that it can obtain its configuration and policy rules.

Troubleshoot License Activation Issues

Some of the most common issues with activating your license are covered in this section.

- Licensing a PA-VM without sufficient memory causes an error similar to the following:

```
Server error : failed key check : Resource check failed.
Memory needed: 6.5GB, allocated memory: 4.8GB
```

To fix this problem, provision the additional memory the license requires, and fetch the license with the command **`request license fetch`**.

If you use any other command it will fail with the following error:

```
Server error : failed to fetch license: Cannot apply a provisioning
 license feature to an already provisioned device.
```

 ©2024 Palo Alto Networks, Inc.

- You see an error that reads **Failed to fetch licenses. Failed to get license info. Please try again later** or a generic communications error message displays.

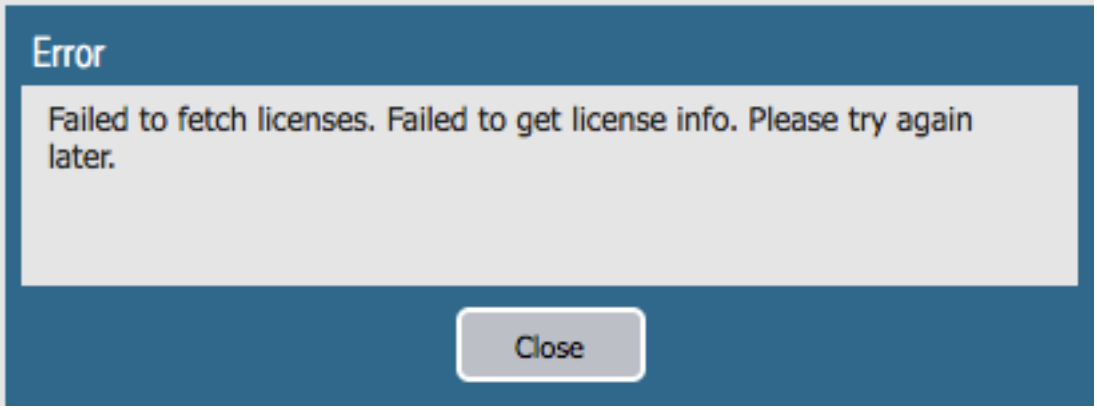

Verify the following:

- Can the firewall route traffic to the Palo Alto Networks server using a service route? By default, the firewall uses the management interface to access the server. If you plan on using a dataplane interface, make sure that you have set up a service route.

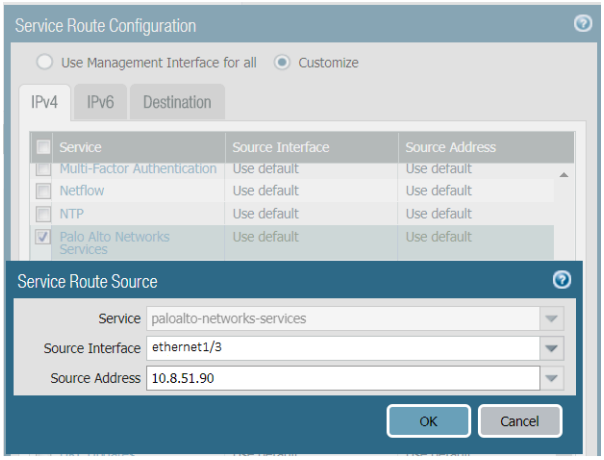

- Is routing over the internet working? SSH into the firewall and ping an publicly accessible IP address such as 4.2.2.2. Be sure to use the source option if you are using a dataplane interface. For example: ping count 3 source 10.0.1.1 host 4.2.2.2.
- Is DNS set up correctly? SSH into the firewall and ping a DNS name such as google.com. For example:

```
warby@warbylan> ping count 3 source 10.0.1.1 host google.com
PING google.com (216.58.195.78) from 10.0.1.1 : 56(84) bytes of data.
64 bytes from sfo07s16-in-f78.1e100.net (216.58.195.78): icmp_seq=1 ttl=55 time=11.6 ms
64 bytes from sfo07s16-in-f78.1e100.net (216.58.195.78): icmp_seq=2 ttl=55 time=11.9 ms
64 bytes from sfo07s16-in-f78.1e100.net (216.58.195.78): icmp_seq=3 ttl=55 time=11.5 ms

--- google.com ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2025ms
rtt min/avg/max/mdev = 11.586/11.721/11.975/0.200 ms
```

 ©2024 Palo Alto Networks, Inc.

- You see an error that reads Invalid Auth Code:

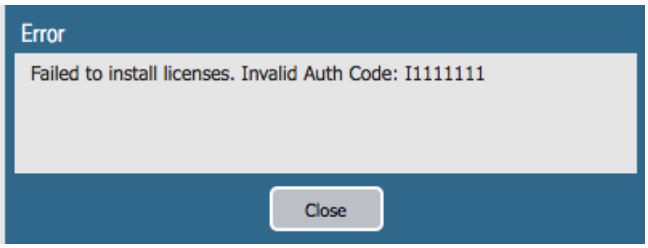

Verify the following:

- You have entered the auth code properly.
- You have registered the auth code to your account on the support portal.
- Your auth code has not reached the maximum provisioning capacity for the VM-Series firewalls.
  - For legacy licenses, log in to the CSP and select **Assets** > **VM-Series Auth-Codes**.
  - For Software NGFW Credits, if you know the deployment profile, log in to the CSP and select **Assets** > **Software NGFW Credits**. Locate your profile, and click **Details**.

    You can also select **Assets** > **Software NGFW Devices** and search by Auth-Code.

Register the VM-Series Firewall

When you purchase a VM-Series firewall, you receive an email that includes an auth code for a capacity license for the VM-Series model, a support entitlement auth code, and one or more auth codes for the subscription licenses. To use the auth code(s), you must register the code to the support account on the Palo Alto Networks Customer Support website. In the case of the VMware integrated NSX solution, the email contains a single authorization code that bundles the capacity license for one or more instances of the VM-Series model, the support entitlement, and one or more subscription licenses.

For the usage-based licenses in public clouds (AWS, Azure, or Google Cloud Platform), you do not receive an auth code. However, in order to activate your premium support entitlement with Palo Alto Networks, you must create a support account and register the VM-Series firewall on the Palo Alto Networks Customer Support website.

Use the instructions in this section to register the capacity auth code or firewall with your support account:

- Register the VM-Series Firewall (Software NGFW Credits)
- Register the VM-Series Firewall (with auth code)
- Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code)

Register the VM-Series Firewall (with auth code)

Complete the following procedure to register your VM-Series firewall with an auth code.

STEP 1 | Log in to the Palo Alto Networks Customer Support website with your account credentials. If you need a new account, see Create a Support Account.

STEP 2 | Select **Assets** > **VM-Series Auth-Codes** > **Add VM-Series Auth-Code**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | In the **Add VM-Series Auth-Code** field, enter the capacity auth code you received by email, and click the checkmark on the far right to save your input. The page will display the list of auth codes registered to your support account.

You can track the number of VM-Series firewalls that have been deployed and the number of licenses that are still available for use against each auth code. When all the available licenses are used, the auth code does not display on the VM-Series Auth-Codes page. To view all the assets that are deployed, select **Assets** > **Devices**.

Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code)

To register usage-based firewalls on the Palo Alto Networks Customer Support Portal (CSP) you can use automatic registration or manual registration. The automatic registration of the usage-based firewalls enables you to seamlessly register the firewall as soon as you launch it and access the site license entitlements associated with your CSP account. For details, see Install a Device Certificate on the VM-Series Firewall.

Use the following workflow to manually register your VM-Series firewalls. Before you begin the manual registration process, log in to the VM-Series firewall and jot down the serial number and the CPU ID (UUID is optional) from the dashboard.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Log in to the Palo Alto Networks Customer Support website, and click **Assets** > **Devices** > **Register Devices**.

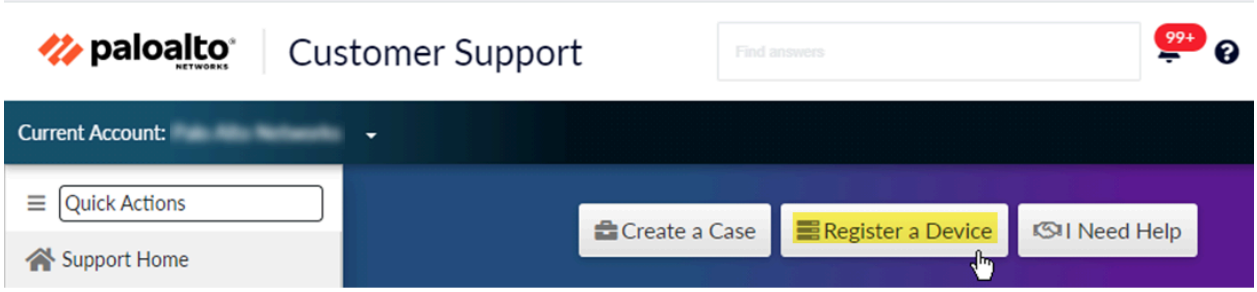

1. Select **Register usage-based VM-Series models (hourly/annual) purchased from public cloud Marketplace or Cloud Security Service Provider (CSSP)**.
2. Select your **Cloud Marketplace** vendor and click **Next**.

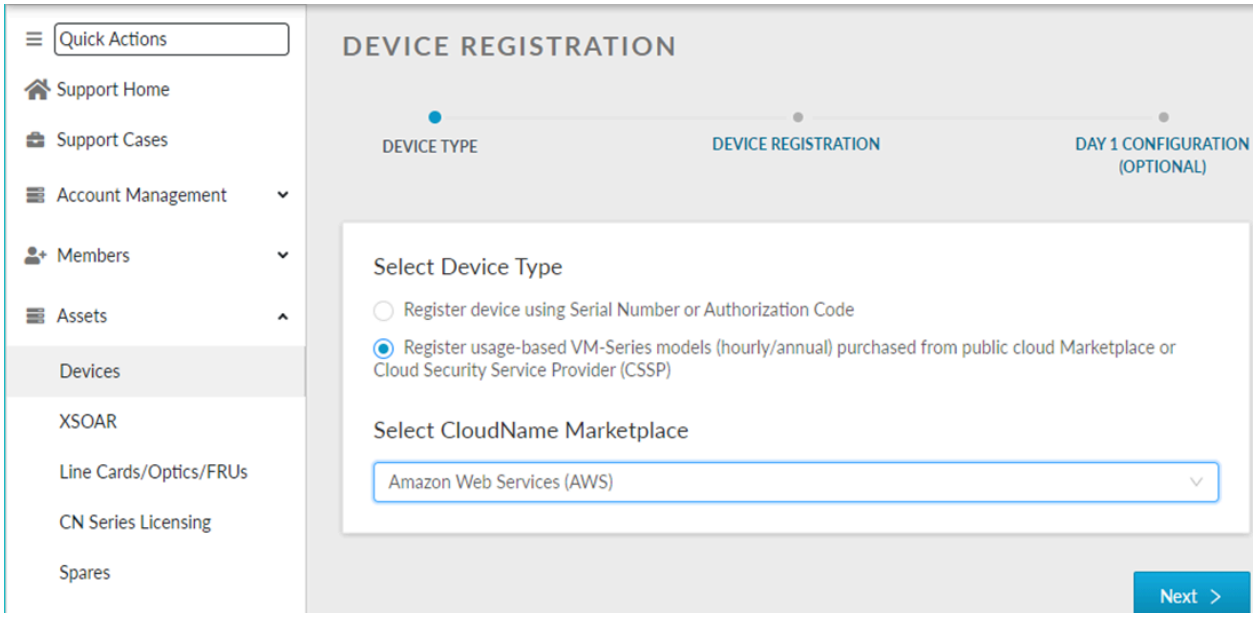

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enter the **Serial #**, the **CPU ID**, and the **UUID** of the VM-Series firewall.

For example, from the Dashboard of the VM-Series firewall on your VM you will see the following information.

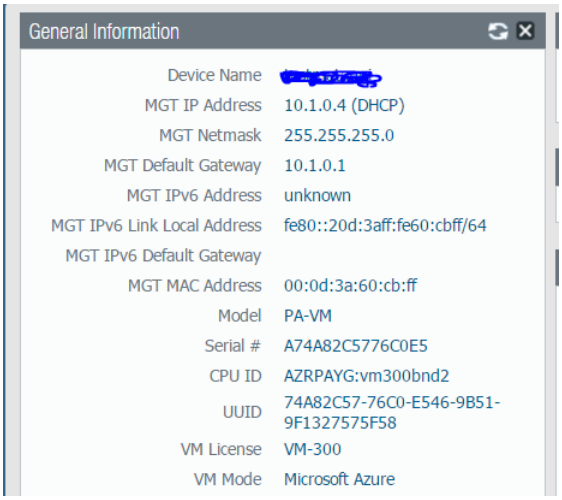

*If you plan to use the firewall offline, please select the **Offline** checkbox and enter the PAN-OS version you plan to use.*

STEP 3 | **Agree and Submit** to accept the EULA and register the firewall.

STEP 4 | Verify that the details on the licenses you purchased are displayed on the CSP **Assets** page.

Install a Device Certificate on the VM-Series Firewall

The firewall requires a device certificate to retrieve the site license entitlements and securely access cloud services such as WildFire, AutoFocus, and Cortex Data Lake. There are two methods for applying a site license to your VM-Series firewall—One-time password and auto-registration PIN. Each password or PIN is generated on the Customer Support Portal (CSP) and unique to your Palo Alto Networks support account. The method you use depends on the license type used to deploy your firewall and if your firewalls are managed by Panorama.

- One Time Password—For VM-Series firewalls that have already been registered with Palo Alto Networks licensing server, you must generate a One-Time Password on the Customer Support Portal and apply it to your VM-Series firewall. Use this method for VM-Series firewalls with a BYOL or ELA license in small-scale, unmanaged deployments and manually-deployed VM-Series firewalls managed by Panorama.
- Registration PIN—This method allows you to apply a site license to your VM-Series firewall at initial startup. Use this method for VM-Series firewalls with usage-based licenses (PAYG), that you bootstrap at launch or with any type of automated deployment, regardless of license type. The auto-registration PIN enables you to automatically register your usage-based firewalls at launch with the CSP and retrieve site licenses.

For the VM-Series firewall on NSX-V and NSX-T, you can add the auto-registration PIN to your service definition configuration so the device certificate is fetched by the firewall upon initial boot up. See the service definition configuration for NSX-V, NSX-T (North-South), and NSX-T (East-West) for more information. If you upgrade previously-deployed firewalls to PAN-OS version that

 ©2024 Palo Alto Networks, Inc.

supports device certificates, you can apply a device certificate to the those firewalls individually using a one-time password.

One-time passwords and auto-registration PINs must be used before they expire. If you do not, you must return to the CSP to generate a new one.

- Enable the firewalls to retrieve the site licenses automatically at launch.

  The firewall requires the device certificate to get the site license entitlements and securely access the cloud services. To retrieve the site licenses when you launch the firewall, in the bootstrap package you must include the authcode in the /license folder, and add the auto registration PIN ID and value in the init-cfg.txt file and place it in the /config folder. Adding the auto registration PIN ID and value also enable auto-registration of the PAYG or usage-based instances on the VM-Series firewall.

  1. Log in to the Customer Support Portal (CSP).
  2. Generate the VM-Series registration PIN.

     Select **Assets** > **Device Certificates** > **Generate Registration PIN**. Save the PIN ID and value. Make sure to launch the firewall before the PIN expires.

     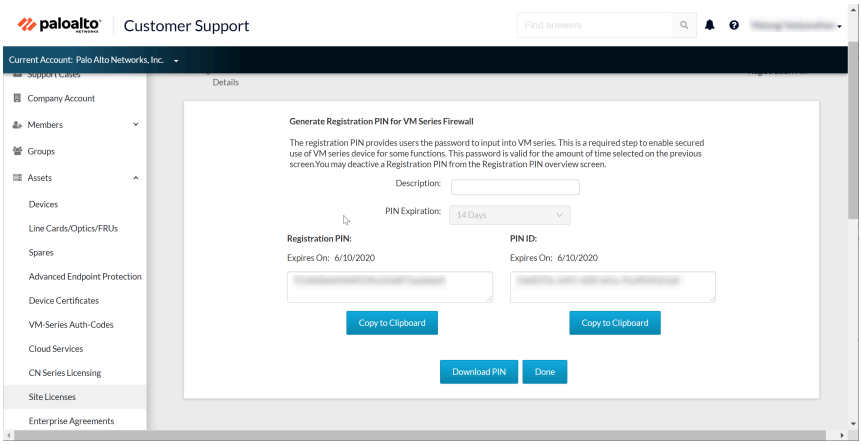

  3. Add the registration PIN ID and value in the init-cfg.txt file.

     In addition to the required parameters, you must include:

     ```
     vm-series-auto-registration-pin-id=
     ```

     ```
     vm-series-auto-registration-pin-value=
     ```

     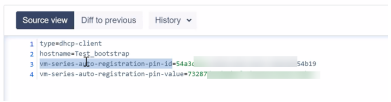

  4. Log in to the firewall and verify that you can see the site license on the firewall.

 ©2024 Palo Alto Networks, Inc.

- Generate the one-time password (OTP) and manually retrieve the device certificate on the firewall.

  The firewall requires this device certificate to get the site license entitlements and securely access the cloud services.

  1. Log in to the Customer Support Portal.

     Register your VM-Series firewall, if you have not already.
  2. Select **Assets** > **Device Certificates** > **Generate OTP**.

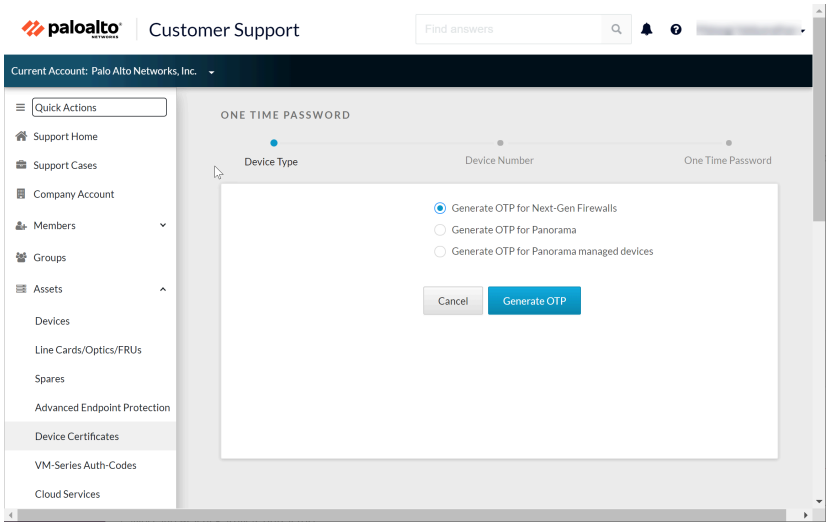

  3. Select your firewall from the list to **Generate OTP**.
  4. Log in to the firewall, and retrieve the device certificate.

     Select **Setup** > **Management** > **Device Certificate** and **Get Certificate**

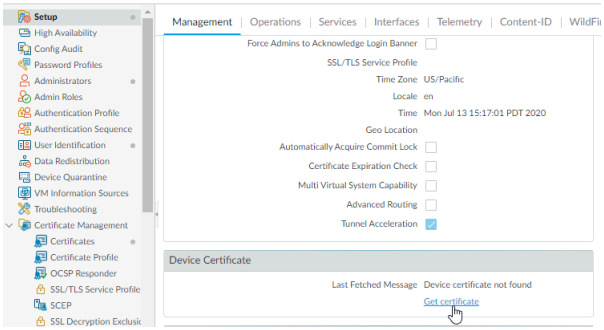

  5. Verify that the device certificate is fetched and that you can see the site license on the firewall.

- If using Panorama to manage the VM-Series firewall, see Install the device certificate on a managed firewall.

Switch Between the BYOL and the PAYG Licenses

The VM-Series firewall cannot be converted between the BYOL and PAYG licensing options. If you have already deployed and configured a VM-Series firewall with the PAYG or BYOL option in AWS, Azure, or Google Cloud Platform, and now want to switch to the other option, use the

 ©2024 Palo Alto Networks, Inc.

following instructions to save and export the configuration on your existing firewall, deploy a new firewall, and then restore the configuration on the new firewall.

STEP 1 | Save a backup of the current configuration file and store it to an external server.

1. Select **Device** > **Setup** > **Operations** and **Export named configuration snapshot**.
2. Select the XML file that contains your running configuration (for example, running-config.xml) and click **OK** to export the configuration file.
3. Save the exported file to a location external to the firewall.

STEP 2 | Deploy a new firewall and register or activate the license, as appropriate.

For a new PAYG instance:

1. In the AWS, Azure, or Google Cloud Platform Marketplace, select the software image for the PAYG licensing bundle you want to deploy.
2. Deploy a new VM-Series firewall in the AWS, Azure, or Google public cloud. See Set Up the VM-Series Firewall on AWS, Set up the VM-Series Firewall on Azure, or Set Up the VM-Series Firewall on Google Cloud Platform.
3. Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code).

For a new BYOL instance:

1. Contact your sales representative or reseller to purchase a BYOL license, and get a BYOL auth code that you can use to license your firewall.
2. Register the VM-Series Firewall (with auth code).
3. Deploy a new VM-Series firewall in the AWS or Azure public cloud. See Set Up the VM-Series Firewall on AWS, Set up the VM-Series Firewall on Azure or Set Up the VM-Series Firewall on Google Cloud Platform.
4. Activate the License for the VM-Series Firewall (Standalone Version).

STEP 3 | On the newly deployed firewall, restore the configuration that you exported.

1. Access the web interface of the newly deployed firewall.
2. Select **Device** > **Setup** > **Operations**, click **Import named configuration snapshot**, Browse to the configuration file on the external host, and click **OK**.
3. Click **Load named configuration snapshot**, select the **Name** of the configuration file you just imported, and click **OK**.
4. Click **Commit** to overwrite the running configuration with the snapshot you just imported.
5. If you have static IP addresses (such as EIP on AWS) attached to your current firewall instance, detach them and attach them to the newly deployed firewall.
6. Verify that the configuration on the new firewall matches the firewall that you are replacing, before you delete the firewall or deactivate the licenses on the replaced firewall.

 ©2024 Palo Alto Networks, Inc.

Switch Between VM-Series Model Licenses

You can switch the license of your currently-deployed VM-Series firewall with the BYOL option. For example, you can move from a subscription bundle to an enterprise license agreement (ELA) and vice versa, without disrupting traffic moving through the firewall. You can also switch the license on an individual firewall or on multiple firewalls simultaneously from Panorama.

Do not use this procedure for switching ELA or perpetual licenses between PAYG and BYOL. See Switch Between the BYOL and the PAYG Licenses for more information.

Complete one of the following procedures to perform one of the following license changes:

- Subscription bundle 1 to subscription bundle 2
- Subscription bundle 1 or 2 to an ELA
- Capacity license to subscription bundle or ELA

Before switching to an ELA license, you must allocate enough tokens to support the number of currently-deployed VM-Series firewalls. See VM-Series Enterprise License Agreement (Multi-Model ELA) for more information about the tokens required for each VM-Series model.

 ©2024 Palo Alto Networks, Inc.

- Switch a license on a standalone firewall.
  1. Register your authorization code.
     - For a subscription bundle, register your new authorization code.
     - For an ELA, activate the ELA authorization code.

     *Do not use the ELA authorization code to activate individual VM-Series firewalls. After registering your ELA, use the VM-Series model authorization codes to activate individual firewalls. You can find these authorization codes on the Customer Support Portal under **Assets** > **VM-Series Auth-Codes**.*
  2. Log in to the VM-Series firewall web interface.
  3. Verify the Palo Alto Networks update server configuration.
     1. Select **Device** > **Setup** > **Services**.
     2. Confirm that **Update Server** is set to updates.paloaltonetworks.com.
     3. Confirm that **Update Server Identity** is selected.
  4. Apply a VM-Series authorization code. A firewall authorization code for an ELA begins with the letter A, as shown below.

     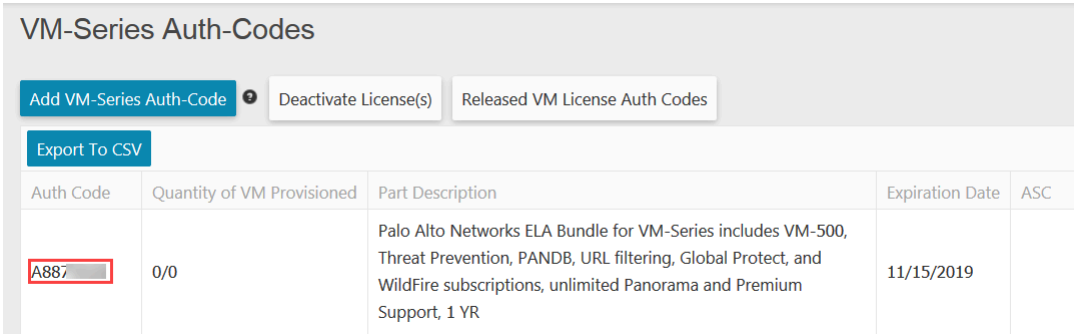

     1. Select **Device** > **Licenses** and select the **Activate feature using authorization code** link.
     2. Enter your VM-Series authorization code.
     3. Click **OK** to confirm the license upgrade.The firewall contacts the Palo Alto Networks update server and consume the tokens required for your firewall based on the VM-Series model.
     4. Verify the license updated successfully by checking the license expiration date.
  5. Repeat this process for each VM-Series firewall in your deployment.

 ©2024 Palo Alto Networks, Inc.

- Switch licenses on managed firewalls using Panorama.
  1. Register your authorization code.
     - For a subscription bundle, register your new authorization code.
     - For an ELA, activate the ELA authorization code.

     *Do not use the ELA authorization code to activate individual VM-Series firewalls. After registering your ELA, use the VM-Series model authorization codes to activate individual firewalls. You can find these authorization codes on the Customer Support Portal under **Assets** > **VM-Series Auth-Codes**.*
  2. Log in to the Panorama web interface.
  3. Verify the Palo Alto Networks update server configuration for the firewalls.
     1. Select **Device** > **Setup** > **Services**.
     2. Confirm that **Update Server** is set to updates.paloaltonetworks.com.
     3. Confirm that **Update Server Identity** is selected.
  4. Apply a VM-Series authorization code. A firewall authorization code for an ELA begins with the letter A, as shown below.

     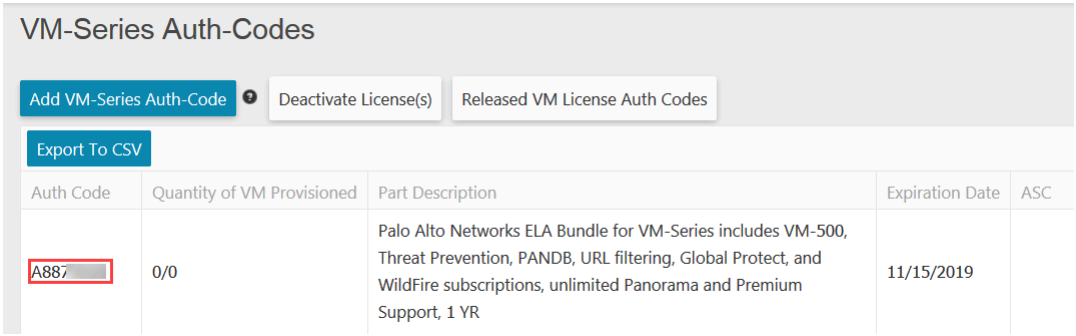

     1. Select **Panorama** > **Device Deployment** > **Licenses** and click **Activate**.
     2. Enter your VM-Series authorization code.
     3. Use the filters to select the managed firewalls to be licensed.
     4. Enter your authorization code in the **Auth Code** column for each firewall.
     5. Click **Activate** to confirm the license upgrade.Panorama contacts the Palo Alto Networks update server and consume the tokens required for your firewalls based on the VM-Series model.

        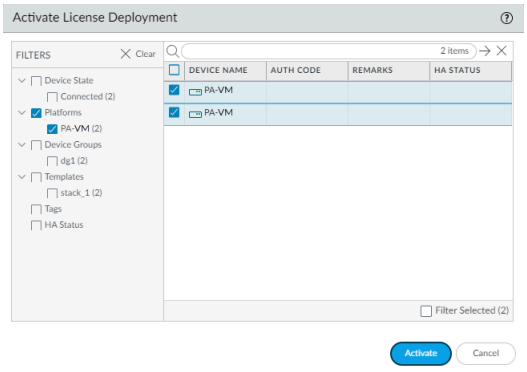

     6. Verify the license updated successfully by checking the license expiration date.

 ©2024 Palo Alto Networks, Inc.

Deactivate License(s)

The license deactivation process enables you to self-manage licenses. Whether you want to remove one or more active licenses or subscriptions attributed to a firewall (hardware-based or VM-Series firewall) or you want to deactivate the VM-Series firewall and unassign all active licenses and subscriptions, begin the deactivation process on the firewall or Panorama (not on the Palo Alto Networks Customer Support web site).

To successfully deactivate a license, you must install a license deactivation API key and enable verification of the update server identity (enabled by default). PAN-OS uses this deactivation API key to authenticate with all update a license services. The deactivation API is key is not required for manual license deactivation, where there is not connectivity between the firewall and license server.

If the firewall/Panorama has internet access and can communicate with the Palo Alto Networks Licensing servers, the license removal process completes automatically with a click of a button. If the firewall/Panorama does not have internet access, you must complete the process manually in a two-step process. In the first step, from the firewall or Panorama, you generate and export a license token file that includes information on the deactivated keys. In the second step, while logged in to the Palo Alto Networks Customer Support website, upload the token file to dissociate the license keys from the firewall.

- Deactivate a Feature License or Subscription Using the CLI
- Deactivate VM

Deactivate a Feature License or Subscription Using the CLI

If you installed a license/subscription on a firewall and need to reassign it to another firewall, you can deactivate the individual license and re-use the same authorization code on another firewall without help from Technical Support. This capability is supported in the CLI only and is supported on both physical and virtual devices running PAN-OS. This procedure is typically used with fixed-model perpetual or ELA licenses.

- Internet Access (Auto Mode)
- No Internet Access (Manual Mode)

Internet Access (Auto Mode)

STEP 1 | Log into the CLI on the firewall.

STEP 2 | View the name of the license key for the feature you want to deactivate.

```
request license deactivate key features
```

STEP 3 | Deactivate the license or subscription.

Use the auto mode to remove the license key.

```
request license deactivate key features <name> mode auto
```

The name is the full name for the license key file. For example:

```
admin@vmPAN2> request license deactivate key features <name>

WildFire_License_2015_01_28_I5820573.key mode auto007200002599
```

 ©2024 Palo Alto Networks, Inc.

```
WildFire License Success
Successfully removed license keys
```

No Internet Access (Manual Mode)

Use manual mode to remove the license key and generate a license token for a model-based license. This procedure assumes you have installed the license API key on your firewall.

STEP 1 | Log into the CLI on the firewall.

STEP 2 | View the name of the license key for the feature you want to deactivate.

request license deactivate key features

STEP 3 | Deactivate the license manually from the command line.

request license deactivate key features <name> mode manual

For example:

admin@PA-VM> request license deactivate key features

```
PAN_DB_URL_Filtering_2015_01_28_I6134084.key mode manual
Successfully removed license keys
dact_lic.01282015.100502.tok
```

The token file uses the format `dact_lic.timestamp.tok`, where the timestamp is in the `dmmyyyy.hrminsec` format.

STEP 4 | Verify that the token file was generated.

show -token-files

STEP 5 | Export the token file.

Enter this command on a single line:

scp export license-token-file to <username@serverIP>

from <token_filename>

For example:

scp export license-token-file to admin@10.1.10.55:/tmp/ from dact_lic.01282015.100502.tok

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Log into the Palo Alto Networks Customer Support portal.

1. Click the **Deactivate License(s)** link on the **Assets** tab.
2. Select **Assets** > **VM-Series Auth-Codes** and select **Deactivate License(s)**.
3. Upload the token file to complete the deactivation.

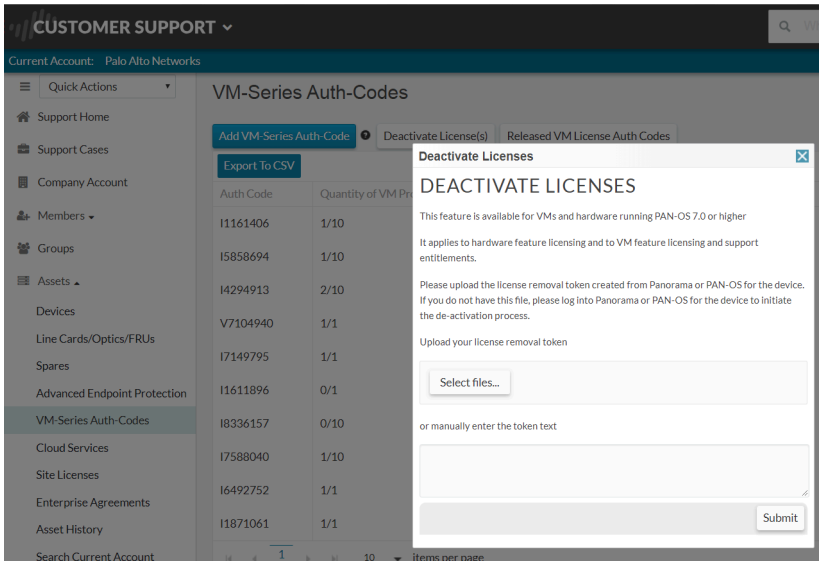

Deactivate VM

When you no longer need a BYOL instance of the VM-Series firewall, you can free up all active licenses (subscription licenses, model-based capacity licenses, and support entitlements) from the web interface, the CLI, or the XML API on the firewall or Panorama. The licenses are credited back to your account and you can use the same authorization codes on a different instance of the VM-Series firewall.

Deactivating a VM removes all the licenses/entitlements and places the VM-Series firewall in an unlicensed state; the firewall will not have a serial number and can support only a minimal number of sessions. Because the configuration on the firewall is left intact, you can re-apply a set of licenses and restore complete functionality on the firewall, if needed.

*Make sure to deactivate licenses **before** you delete the VM-Series firewall. If you delete the firewall before deactivating the licenses, you have two options:*

***Managed by Panorama**—Deactivate the license from Panorama.*

***Not managed by Panorama**—Contact Palo Alto Networks Customer Support for deactivation assistance.*

- Deactivate the VM from the Firewall
- Deactivate the VM from Panorama

Deactivate the VM from the Firewall

Complete the following process to deactivate the VM license from the firewall.

STEP 1 | Log into the web interface and select **Device** > **Licenses**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | In the License Management section, select **Deactivate VM**.

You can only see this option on a VM. It is not there on a physical firewall.

STEP 3 | Verify the list of licenses/entitlements to be deactivated on the firewall.

STEP 4 | Pick one of the following options to start deactivating the VM:

- (Internet access to the Palo Alto Networks Licensing server) Select **Continue**.

  You are prompted to reboot the firewall; on reboot the licenses are deactivated.
- (No internet)—Select **Complete Manually**.

  Click the **Export license token** link to save the token file to your local computer. Here is a sample token filename: 20150128_1307_dact_lic.01282015.130737.tok

  You are prompted to reboot the firewall; on reboot the licenses are deactivated.

STEP 5 | (Manual Process—no internet) Use the token file to register the changes with the Licensing server:

1. Log into the Palo Alto Networks Customer Support website.
2. Select **Assets** > **VM-Series Auth-Codes** > **Deactivate License(s)**.
3. While logged in to the Palo Alto Networks Customer Support website, upload the token file to complete the deactivation.

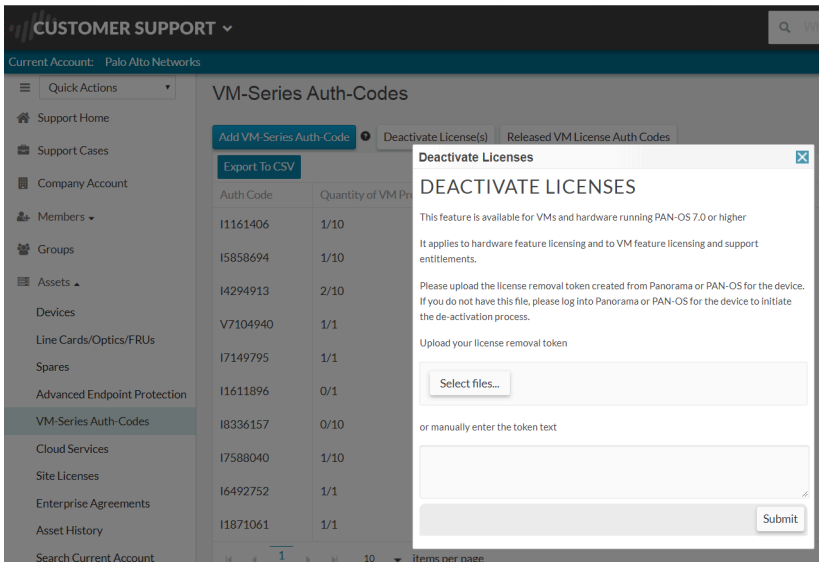

Deactivate the VM from Panorama

Complete the following process to deactivate a VM license from Panorama.

STEP 1 | Log in to the Panorama web interface and select **Panorama** > **Device Deployment** > **Licenses**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | **Deactivate VMs** and select the VM-Series firewall that you want to deactivate.

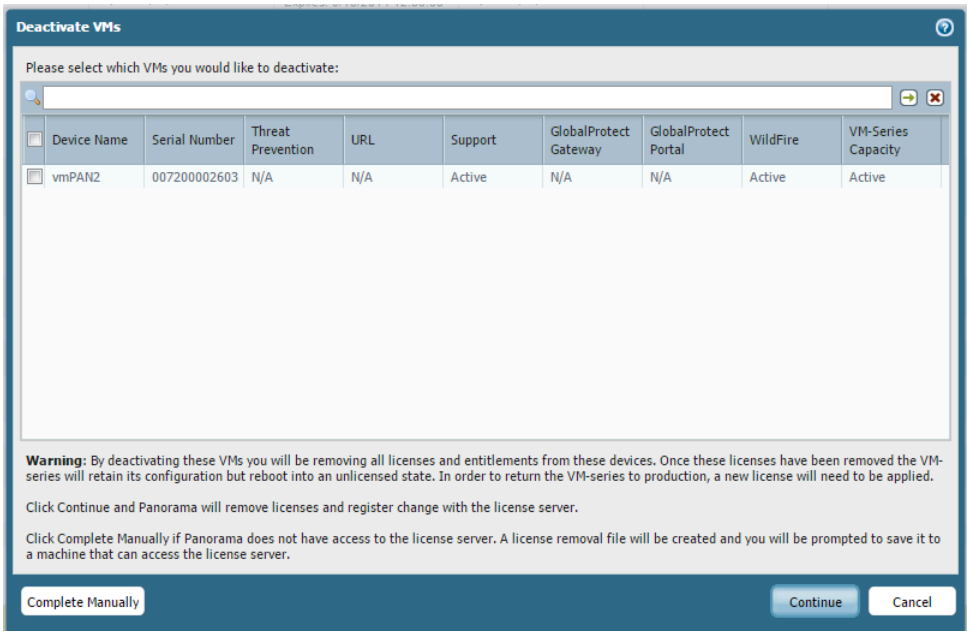

STEP 3 | Pick one of the following options to deactivate the VM:

- **Continue**—If Panorama can communicate directly with the Palo Alto Networks Licensing servers and can register the changes. To verify that the licenses have been deactivated on

 ©2024 Palo Alto Networks, Inc.

the firewall, select **Refresh on Panorama** > **Device Deployment** > **Licenses**. The firewall is automatically rebooted.

- **Complete Manually**—If Panorama does not have internet access, Panorama generates a token file.

  Click the **Export license token** link to save the token file to your local computer. Here is a sample token filename: 20150128_1307_dact_lic.01282015.130737.tok

  The successful completion message is displayed on-screen, and the firewall is automatically rebooted.

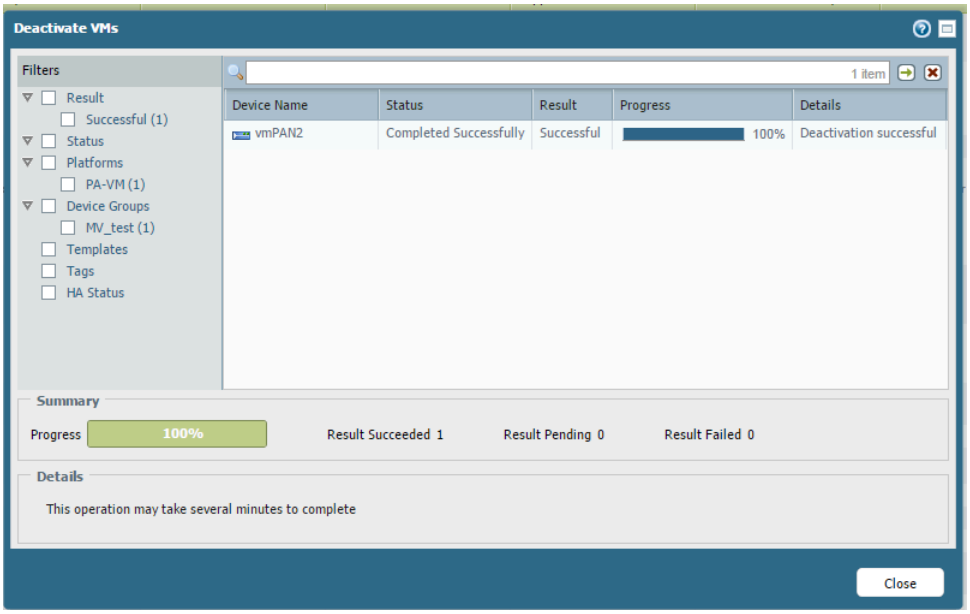

STEP 4 | (Manual process only —no internet) Use the token file to register the changes with the licensing server.

1. Log into the Palo Alto Networks Customer Support website.
2. Select **Assets** > **VM-Series Auth-Codes** > **Deactivate License(s)**.
3. Upload the token file to complete the deactivation.

STEP 5 | Remove the deactivated VM-Series firewall as a managed device on Panorama.

1. Select **Panorama** > **Managed Devices**.
2. Select the firewall that you deactivated from the list of managed devices, and click **Delete**.

   Instead of deleting the firewalls, if you prefer, you can create a separate device group and assign the deactivated VM-Series firewalls to this device group.

Renew VM-Series Firewall License Bundles

When your VM-Series firewall bundle licenses are due for renewal, you can log in to the Palo Alto Networks Customer Support Portal and adjust the license quantity to meet your deployment needs. At renewal, you can review your usage trends and based your future needs, pick from the following options:

 ©2024 Palo Alto Networks, Inc.

- **Renew**—You can opt to renew all licenses as is, or to increase or decrease the licensed quantity. If you decrease the number of licenses you need, you must opt to get a basic bundle for the firewalls you are not renewing, otherwise you will forfeit the portion that you do not renew. If you increase the license quantity, the addition is added to your existing auth code.
- **Change to Basic Bundle**—If you have a VM-Series bundle 1 or a bundle 2 license that includes subscriptions, you can change to a basic bundle that includes a perpetual capacity license and support entitlement. When you switch to the basic bundle, you retain the VM-Series firewall model that you had previously purchased. All firewalls that are currently deployed and are associated with the existing auth code will continue to function, and the support entitlement will have a new expiration date. For any unprovisioned firewalls, you'll receive a new auth code that you can use to deploy new instances.
- **Forfeit**—Relinquish the licenses that you no longer need. If you have deployed the firewalls that you don't want to renew, you need to select the serial number of the instances for which you want to discontinue renewals. You can continue to use these firewall instances with the software and content versions that are currently installed, but your subscriptions and support entitlements are no longer valid. And to forfeit the license of VM-Series firewalls that you have not provisioned, just select the quantity that you want to forfeit.

STEP 1 | Log in to the Palo Alto Networks Customer Support Portal with your account credentials.

STEP 2 | Select **Assets** > **VM-Series Auth-Codes** and find the auth code you want to renew.

The **Renew** option displays for auth codes that are eligible for renewal.

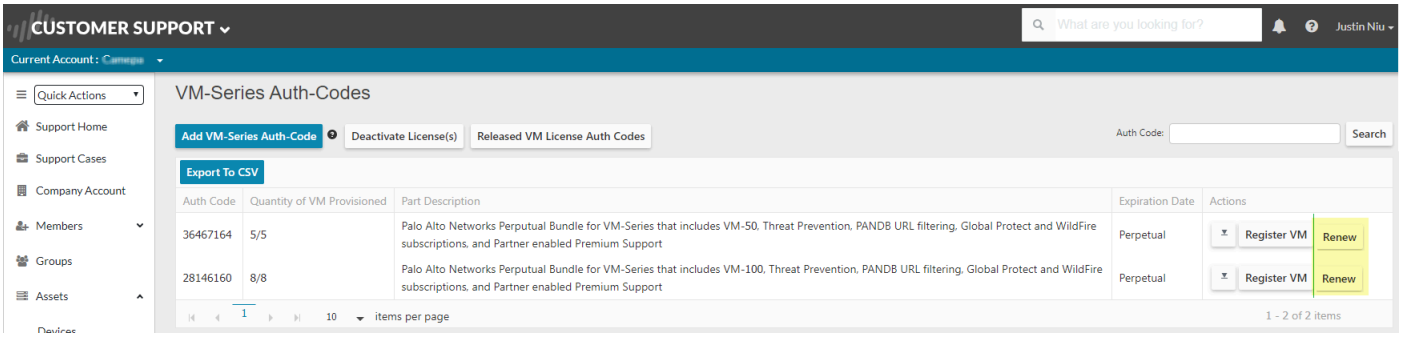

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Click the **Renew** link to select the serial numbers to **Renew**, **Change to Basic Bundle**, or **Forfeit**.

If you have provisioned the firewall, select the appropriate option in the row that corresponds to the Serial Number. If you have unprovisioned instances of the firewall, select the quantity for each renewal option you choose under **Unprovisioned VM Renewal Settings**.

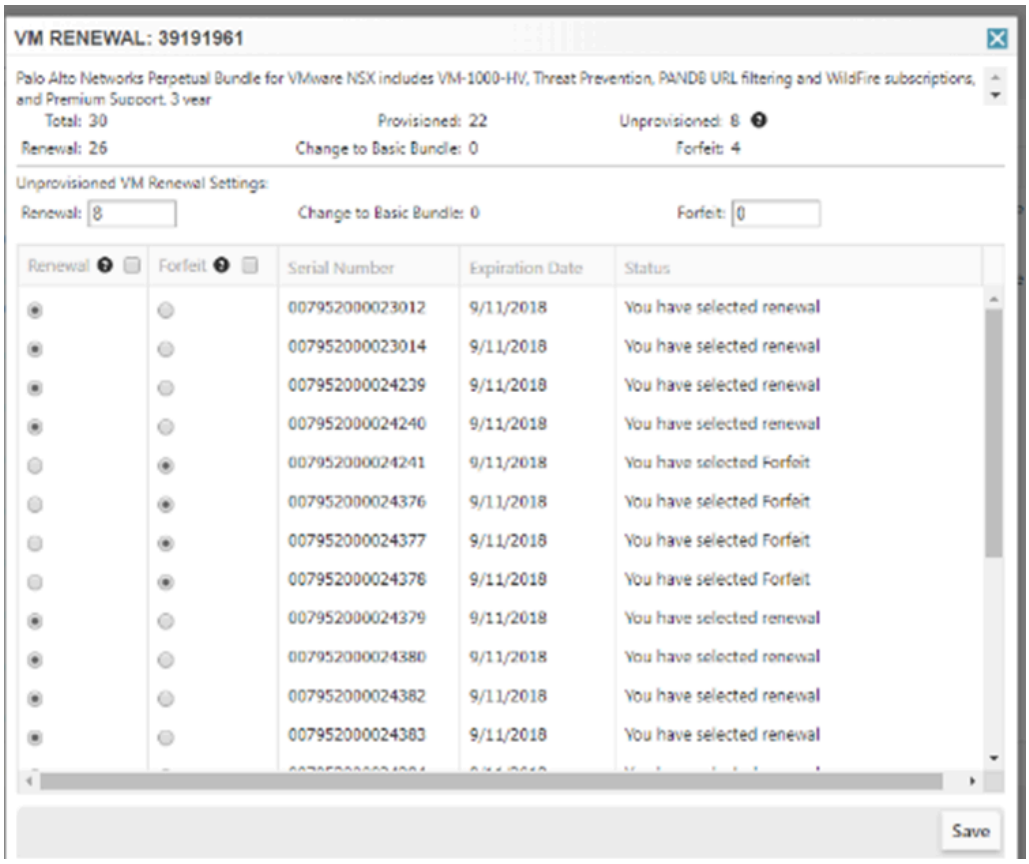

STEP 4 | **Save** your changes.

You will receive an onscreen confirmation that your changes are submitted for processing. After submitting your changes, if you select Renew again, you can view the status of your request against each serial number. If renewal processing has started, and you need to make additional revisions, you will be unable to save changes. For assistance, you can contact the renewals team at renewals@paloaltonetworks.com.

 ©2024 Palo Alto Networks, Inc.

What Happens When Licenses Expire?

Palo Alto Networks VM-Series firewall licenses and subscriptions provide the firewall with added functionality and/or access to a Palo Alto Networks cloud-delivered service. When a license is within 30 days of expiration, a warning message displays in the system log daily until the subscription is renewed or expires. Upon license expiration, some subscriptions continue to function in a limited capacity, and others stop operating completely. Here you can find out what happens when each subscription expires.

The precise moment of license expiry is at the beginning of the following day at 12:00 AM (GMT). For example, if your license is scheduled to end on 1/20 you will have functionality for the remainder of that day. At the start of the new day on 1/21 at 12:00 AM (GMT), the license will expire. All license-related functions operate on Greenwich Mean Time (GMT), regardless of the configured time zone on the firewall.

License	Expiry Behavior
VM-Series	**You can still:** • You can continue to configure and use the firewall(s) you deployed prior to the license expiring with no change in session capacity. The firewall(s) will not reboot automatically and cause a disruption in traffic. However, if the firewall reboots for any reason, the firewall enters an unlicensed state. While unlicensed, the number of supported sessions is limited to 1200. No other management plane features or configuration options are restricted.
Threat Prevention	Alerts appear in the System Log indicating that the license has expired. **You can still:** • Use signatures that were installed at the time the license expired, unless you install a new Applications-only content update either manually or as part of an automatic schedule. If you do, the update will delete your existing threat signatures and you will no longer receive protection against them. • Use and modify Custom App-ID™ and threat signatures. **You can no longer:** • Install new signatures. • Roll signatures back to previous versions.
DNS Security	**You can still:**

 ©2024 Palo Alto Networks, Inc.

License	Expiry Behavior
	• Use local DNS signatures if you have an active Threat Prevention license. **You can no longer:** • Get new DNS signatures.
Advanced URL Filtering / URL Filtering	**You can still:** • Enforce policy using custom URL categories. • Enforce policy using PAN-DB categories that were in your local cache when the license expired. **You can no longer:** • Get updates to cached PAN-DB categories. • Connect to the PAN-DB URL filtering database. • Get PAN-DB categories of uncached URLs. • Analyze URL requests in real-time using advanced URL filtering.
WildFire	**You can still:** • Forward PEs for analysis. • Get signature updates every 24-48 hours if you have an active Threat Prevention subscription. **You can no longer:** • Get five-minute updates through the WildFire public and private clouds. • Forward advanced file types such as APKs, Flash files, PDFs, Microsoft Office files, Java Applets, Java files (.jar and .class), and HTTP/HTTPS email links contained in SMTP and POP3 email messages. • Use the WildFire API. • Use the WildFire appliance to host a WildFire private cloud or a WildFire hybrid cloud.
AutoFocus	**You can still:** • Use an external dynamic list with AutoFocus data for a grace period of three months. **You can no longer:** • Access the AutoFocus portal. • View the AutoFocus Intelligence Summary for Monitor log or ACC artifacts.

 ©2024 Palo Alto Networks, Inc.

License	Expiry Behavior
Cortex Data Lake	**You can still:** • Store log data for a 30-day grace period, after which it is deleted. • Forward logs to Cortex Data Lake until the end of the 30-day grace period.
GlobalProtect	**You can still:** • Use the app for endpoints running Windows and macOS. • Configure single or multiple internal/external gateways. **You can no longer:** • Access the Linux OS app and mobile app for iOS, Android, Chrome OS, and Windows 10 UWP. • Use IPv6 for external gateways. • Run HIP checks. • Use Clientless VPN. • Enforce split tunneling based on destination domain, client process, and video streaming application.
Support	**You can no longer:** • Receive software updates. • Download VM images. • Benefit from technical support.

 ©2024 Palo Alto Networks, Inc.

Licenses for Cloud Security Service Providers (CSSPs)

The Palo Alto Networks CSSP partners program allows service providers to provide security as a service or as a hosted application to their end customers. The license offerings that Palo Alto Networks provides for authorized Cloud Security Service Provider (CSSP) partners are different from the offerings for enterprise users.

For CSSP partners, Palo Alto Networks supports a usage-based model for the VM-Series firewalls bundled with subscriptions and support. CSSP partners can combine a term-based capacity license for the VM-Series Models with a choice of subscription licenses for Threat Prevention, URL Filtering, AutoFocus, GlobalProtect, and WildFire, and support entitlements that provide access to technical support and software updates. If you plan on deploying the firewalls in an HA configuration, you can purchase the cost-effective high availability option.

- Get the Auth Codes for CSSP License Packages
- Register the VM-Series Firewall with a CSSP Auth Code
- Add End-Customer Information for a Registered VM-Series Firewall

Get the Auth Codes for CSSP License Packages

To be a CSSP Partner, you have to enroll in the Palo Alto Networks CSSP partners program. For information on enrolling in the CSSP program, contact your Palo Alto Networks Channel Business Manager. If you are enrolled, the Palo Alto Network Support portal provides tools that allow you to select a license package, track license usage, and apply license entitlements.

A license package is a combination of the following options:

- Usage term—The pay-per-use options are hourly, monthly, 1-year, and 3-years.
- VM-Series firewall model—The VM-100, VM-200, VM-300, and VM-1000-HV that give you the model number and the capacities associated with each model.
- Subscription bundle—The three options are basic, bundle 1, and bundle 2. The basic option does not include any subscriptions; bundle 1 has the Threat Prevention license that includes IPS, AV, malware prevention; bundle 2 has the Threat Prevention (includes IPS, AV, malware prevention), DNS Security, GlobalProtect, WildFire, and PAN-DB URL Filtering licenses.
- Level of support—Premium support or backline support.
- Redundant firewalls—The option are either high availability (HA) or without HA. This option is a cost-effective option if you plan to deploy a pair of redundant firewalls.

The offering PAN-VM-300-SP-PREM-BND1-YU, for example, is a one-year term package that includes the VM-300 with premium support and the subscription bundle 1. Each package supports up to a maximum of 10,000 instances of the VM-Series firewall.

After you select your license package, you receive an email with your auth code; the fulfillment process can take up to 48 hours.

STEP 1 | Log in to the Palo Alto Networks Customer Support website with your account credentials. If you need a new account, see Create a Support Account.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Select **CSSP** > **Order History**, to view the list of auth codes registered to your support account.

As you deploy firewalls, you must register each instance of the firewall against an auth code.

Register the VM-Series Firewall with a CSSP Auth Code

To activate the license on your VM-Series firewall, you must have deployed the VM-Series firewall and completed initial configuration. As a CSSP partner, you can choose from the following options to register a firewall:

- API—Use the Licensing API if you have a custom script or an orchestration service. With this option, the firewall does not need direct internet access.
- Bootstrap—Use this option to automatically configure the firewall and license it on first boot. See Bootstrap the VM-Series Firewall.
- Firewall web interface—You can Activate the License for the VM-Series Firewall (Standalone Version) using the firewall web interface. This workflow is valid for firewalls with or without internet access.
- Customer Support Portal—Use this option to manually register the firewall on the Palo Alto Networks Customer Support portal, as shown below.

STEP 1 | Log in to the Palo Alto Networks Customer Support website with your account credentials. If you need a new account, see Create a Support Account.

STEP 2 | Select **CSSP** > **Order History**, to view the list of auth codes registered to your support account.

STEP 3 | Select **CSSP** > **VM Provisioning Auth Codes**, select an **Authorization Code** and click **Register VM**.

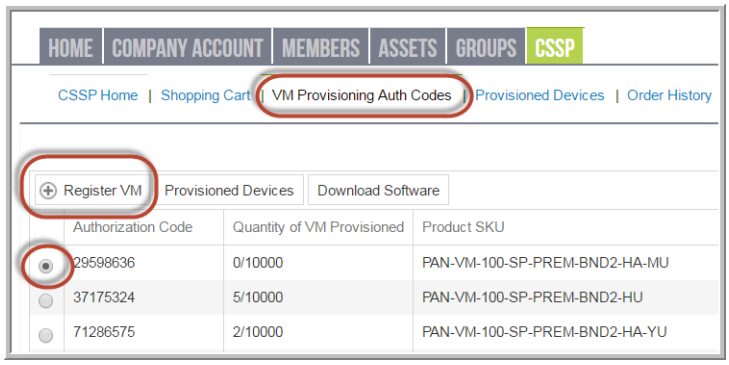

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enter the **UUID** and **CPUID** of the VM instance and click **Submit**. The portal will generate a serial number for the firewall.

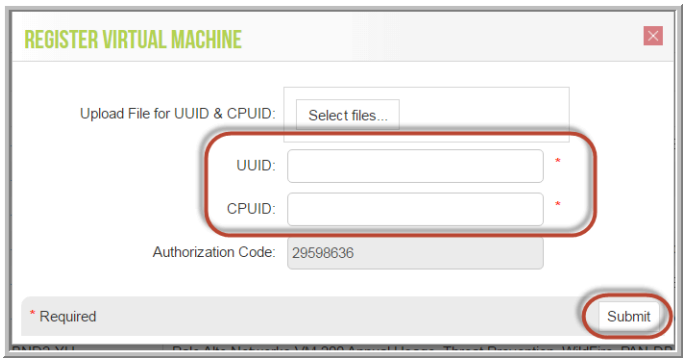

*You can track the number of VM-Series firewalls that have been deployed and the number of licenses that are still available for use against each auth code. To view all the total number of firewalls registered against a specific auth code, select **CSSP** > **VM Provisioning Auth Codes**, then select an **Authorization Code** and click **Provisioned Devices**.*

Add End-Customer Information for a Registered VM-Series Firewall

For CSSP licensees, after you register the firewall, you can use either the Palo Alto Networks Support portal or the Licensing API to link the serial number of the VM-Series firewall with the customer for whom you provisioned the firewall.

- Add End-Customer Information for a Registered VM-Series Firewall (Customer Support Portal). The Support portal authenticates with user name and password.
- Add End-Customer Information for a Registered VM-Series Firewall (API). The API authenticates using the Licensing API key.

Add End-Customer Information for a Registered VM-Series Firewall (Customer Support Portal)

Complete the following procedure to add end-customer information for a registered firewall through the Customer Support Portal.

STEP 1 | Log in to the Palo Alto Networks Customer Support website with your account credentials.

STEP 2 | Select **CSSP** > **Provisioned Devices**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Select the **Serial Number** and click **Add End User Info**.

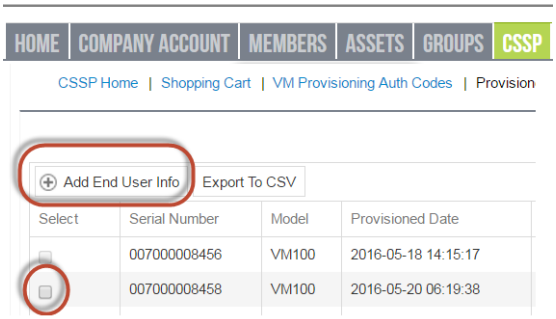

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enter the **Account Information** for the customer as follows.

- Customer Reference Id: Required
- Company Name: Required
- DNB #: Data Universal Numbering System (D-U-N-S) number
- Contact Email: Required, end-user email address
- Contact Phone Number: End-user phone number
- Address: Required, end-user address
- Country: Required, ISO 2-letter country code
- City: Required, end-user city name
- Region/State: Required; for the United States and Canada, you must enter an ISO 2-letter subdivision code; for all other countries, any text string is valid
- Postal Code: Required, end-user postal code
- Company Website: End-user website URL
- Industry: End-user industry type, such as networking or consultancy

Click **Submit** to save the details.

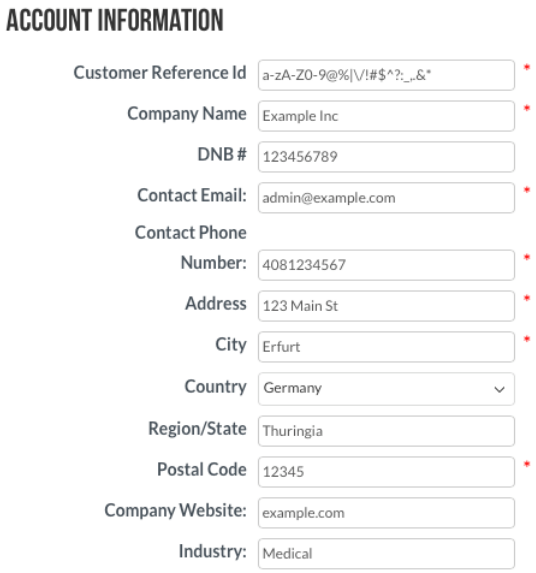

*After you add account information, you can find all firewalls registered to a customer. In **Search Existing End User**, enter the customer ID or customer name and click **Search** to find all firewalls provisioned for the customer.*

Add End-Customer Information for a Registered VM-Series Firewall (API)

The URL for accessing the API is https://api.paloaltonetworks.com/api/license/ReportEndUserInfo.

An API request must use the HTTP POST method, and you must include HTTP requests headers that include the API key and specify the content type as JSON. API responses are in JSON format.

STEP 1 | Get your Licensing API key.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Use the ReportEndUserInfo API to add end-user information for a VM-Series Firewall that is registered to a CSSP.

URL: **`https://api.paloaltonetworks.com/api/license/ReportEndUserInfo`**

Headers:

- Content-Type: application/json
- apiKey: *API Key*

Parameters:

- `SerialNumbers`: Required, provide at least one valid firewall serial number
- `CustomerReferenceId`: Required
- `CompanyName`: Required, end-user company name
- `DnBNumber`: Data Universal Numbering System (D-U-N-S) number
- `PhoneNumber`: End-user phone number
- `EndUserContactEmail`: Required, end-user email address
- `Address`: Required, end-user address
- `Country`: Required, ISO 2-letter country code
- `City`: Required, end-user city name
- `Region/State`: Required; for the United States and Canada, you must enter an ISO 2-letter subdivision code; for all other countries, any alpha string is valid
- `PostalCode`: Required, end-user postal code
- `Industry`: End-user industry type, such as networking or consultancy
- `WebSite`: End-user website URL
- `CreatedBy`: System or person submitting this information

Sample request to add end-user information for a registered VM-Series firewall using Curl:

```
curl -X POST "http://api.paloaltonetworks.com/api/license/
ReportEndUserInfo" \-H "Content-Type: application/json" \-
H "apikey: your_key_here" \--data-raw '{ "SerialNumbers":
 ["0001A101234"], "CustomerAccountId": 12345, "CompanyName":
 "ExampleInc", "DnBNumber": "123456789", "Address": "123 Main St",
 "City": "Sunnydale", "Region": "CA", "State": "CA", "Country":
 "US", "PostalCode": "12345", "Industry": "Medical", "PhoneNumber":
 "4081234567", "WebSite": "example.com", "EndUserContactEmail":
 "admin@example.com", "CreatedBy": "Jane Doe"}'
```

Sample API response:

```
"{"Message": "End User Information Updated Successfully"}"
```

If you receive an error, see Licensing API Error Codes.

 ©2024 Palo Alto Networks, Inc.

Set Up a VM-Series Firewall on an ESXi Server

The VM-Series firewall is distributed in the Open Virtualization Alliance (OVA) format, which is a standard method of packaging and deploying virtual machines. You can install this solution on any x86 device that is capable of running VMware ESXi.

In order to deploy a VM-Series firewall you must be familiar with VMware and vSphere, including vSphere networking, ESXi host setup and configuration, and virtual machine guest deployment.

If you want to automate the process of deploying a VM-Series firewall, you can create a gold standard template with the optimal configuration and policies, then use the vSphere API and the PAN-OS XML API to rapidly deploy new VM-Series firewalls in your network.

See the following topics for information:

- Supported Deployments on VMware vSphere Hypervisor (ESXi)
- VM-Series on ESXi System Requirements and Limitations
- Install a VM-Series firewall on VMware vSphere Hypervisor (ESXi)
- VM Monitoring on vCenter
- Troubleshoot ESXi Deployments
- Performance Tuning of the VM-Series for ESXi

Supported Deployments on VMware vSphere Hypervisor (ESXi)

You can deploy one or more instances of the VM-Series firewall on the ESXi server. Where you place the VM-Series firewall on the network depends on your topology. Choose from the following options (for environments that are not using VMware NSX):

- **One VM-Series firewall per ESXi host**—Every VM server on the ESXi host passes through the firewall before exiting the host for the physical network. VM servers attach to the firewall via virtual standard switches. The guest servers have no other network connectivity, therefore the firewall has visibility and control over all traffic leaving the ESXi host. One variation of this use case is to also require all traffic to flow through the firewall, including server to server (east-west) traffic on the same ESXi host.
- **One VM-Series firewall per virtual network**—Deploy a VM-Series firewall for every virtual network. If you have designed your network such that one or more ESXi hosts has a group of virtual machines that belong to the internal network, a group that belongs to the external network, and a group that belongs to the DMZ, you can deploy a VM-Series firewall to safeguard the servers in each group. If a group or virtual network does not share a virtual switch or port group with any other virtual network, it is completely isolated from all other virtual networks within or across the host(s). Because there is no other physical or virtual path to any other network, the servers on each virtual network must use the firewall to talk to any other network. The firewall has visibility and control over all traffic leaving the virtual (standard or distributed) switch attached to each virtual network.
- **Hybrid environment**—Both physical and virtual hosts are used. The VM-Series firewall can replace a physical firewall appliance in a traditional aggregation location. A hybrid environment achieves the benefits of a common server platform for all devices, and unlinks hardware and software upgrade dependencies.

Continue with VM-Series on ESXi System Requirements and Limitations and Install a VM-Series firewall on VMware vSphere Hypervisor (ESXi).

 ©2024 Palo Alto Networks, Inc.

VM-Series on ESXi System Requirements and Limitations

This section lists requirements and limitations for the VM-Series firewall on VMware vSphere Hypervisor (ESXi). To deploy the VM-Series firewall, see Install a VM-Series firewall on VMware vSphere Hypervisor (ESXi).

- VM-Series on ESXi System Requirements
- VM-Series on ESXi System Limitations

VM-Series on ESXi System Requirements

You can create and deploy multiple instances of the VM-Series firewall on an ESXi server. Because each instance of the firewall requires a minimum resource allocation—number of CPUs, memory and disk space—on the ESXi server, make sure to conform to the specifications below to ensure optimal performance.

The VM-Series firewall has the following requirements:

- The host CPU must be an x86-based Intel or AMD CPU with virtualization extension.
- See the Compatibility Matrix for supported versions of ESXi. The support for the vmx version is based on the OVA that you use to deploy the VM-Series firewall, and you cannot modify this version. Upgrading or downgrading the VM-Series software version does not change the vmx version that was enabled at launch.
- See VM-Series System Requirements for the minimum hardware requirements for your VM-Series model.
- Minimum of two network interfaces (vNICs). One is a dedicated vNIC for the management interface and one is for the data interface. You can then add up to eight more vNICs for data traffic. For additional interfaces, use VLAN Guest Tagging (VGT) on the ESXi server or configure subinterfaces on the firewall.

  Hypervisor assigned MAC address are enabled by default. vSphere assigns a unique vNIC MAC address to each dataplane interface of the VM-Series firewall. If you disable hypervisor assigned MAC addresses, the VM-Series firewall assigns each interface a MAC address from its own pool. Because this causes the MAC addresses on each interface to differ, you must enable promiscuous mode on the port group of the virtual switch to which the firewall's dataplane interfaces are attached; this allows the firewall to receive frames (see Provision the VM-Series Firewall on an ESXi Server). If neither promiscuous mode nor hypervisor assigned MAC address is enabled, the firewall does not receive any traffic. This is because vSphere does not forward frames to a virtual machine when the frame's destination MAC address and the vNIC MAC address do not match.
- Data Plane Development Kit (DPDK) is enabled by default on VM-Series firewalls on ESXi. For more information about DPDK, see Enable DPDK on ESXi.

 ©2024 Palo Alto Networks, Inc.

- To achieve the best performance out of the VM-Series firewall, you can make the following adjustments to the host before deploying the VM-Series firewall. See Performance Tuning of the VM-Series for ESXi for more information.
  - **Enable DPDK**. DPDK allows the host to process packets faster by bypassing the Linux kernel. Instead, interactions with the NIC are performed using drivers and the DPDK libraries.
  - **Enable SR-IOV**. Single root I/O virtualization (SR-IOV) allows a single PCIe physical device under a single root port to appear to be multiple separate physical devices to the hypervisor or guest.

    Do not configure a vSwitch on the physical port on which you enable SR-IOV. To communicate with the host or other virtual machines on the network, the VM-Series firewall must have exclusive access to the physical port and associated virtual functions (VFs) on that interface.
  - **Enable multi-queue support for NICs**. Multi-queue allows network performance to scale with the number of vCPUs and allows for parallel packet processing by creating multiple TX and RX queues.

VM-Series on ESXi System Limitations

The VM-Series firewall functionality is very similar to the Palo Alto Networks hardware firewalls, but with the following limitations:

- Do not use the VMware snapshots functionality on the VM-Series on ESXi. Snapshots can impact performance and result in intermittent and inconsistent packet loss.See the VMware best practice recommendation for using snapshots.

  If you need configuration backups, use Panorama, or from the firewall, use **Export named configuration snapshot** (Device > Set up > Operations). Using **Export named configuration snapshot** exports the firewall's active configuration (`running-config.xml`) and allows you to save it to any network location.
- Dedicated CPU cores are recommended.
- High Availability (HA) Link Monitoring is not supported on VM-Series firewalls on ESXi. Use Path Monitoring to verify connectivity to a target IP address or to the next hop IP address.
- Up to 10 total ports can be configured; this is a VMware limitation. One port is used for management traffic and up to 9 can be used for data traffic.
- Only the vmxnet3 driver is supported.
- Virtual systems are not supported.
- vMotion of the VM-Series firewall is supported on vSphere 6.5, 6.7, and 7.0 if the ESXi hosts have homogeneous CPU configuration. PAN-OS 9.1.6 and later is required to Use vMotion to Move the VM-Series Firewall Between Hosts installed on vSphere 6.5 or 6.7.
- Forged transmit and promiscuous mode must be enabled on the ESXi vSwitch port groups connected to Layer 2 and vwire interfaces on the VM-Series firewall.
- To use PCI devices with the VM-Series firewall on ESXi, memory mapped I/O (MMIO) must be below 4GB. You can disable MMIO above 4GB in your server's BIOS. This is an ESXi limitation.
- When using ESXi 7.0, interfaces do not come up when attaching VFs to virtual machines with PCI device passthrough.

 ©2024 Palo Alto Networks, Inc.

Install a VM-Series firewall on VMware vSphere Hypervisor (ESXi)

To install a VM-Series firewall you must have access to the Open Virtualization Alliance format (OVA) template. Use the auth code you received in your order fulfillment email to register your VM-Series firewall and download the OVA template. The OVA template is a zip archive that contains three types of files:

- .mf: OVF manifest file that contains the SHA-1 digests of individual files in the package
- .ovf: OVF descriptor file that contains all metadata for the package and its contents
- .vmdk: Virtual disk image file that contains the virtualized version of the firewall

Complete the following tasks to install and configure the VM-Series firewall on ESXi.

- Plan the Interfaces for the VM-Series for ESXi
- Provision the VM-Series Firewall on an ESXi Server
- Perform Initial Configuration on the VM-Series on ESXi
- (Optional) Add Additional Disk Space to the VM-Series Firewall
- Use VMware Tools on the VM-Series Firewall on ESXi and vCloud Air
- Use vMotion to Move the VM-Series Firewall Between Hosts
- Use the VM-Series CLI to Swap the Management Interface on ESXi

Plan the Interfaces for the VM-Series for ESXi

By planning the mapping of VM-Series Firewall vNICs and interfaces, you can avoid reboots and configuration issues. The following table describes the default mapping between VMware vNICs and VM-Series interfaces when all 10 vNICs are enabled on ESXi.

VMware vNIC	VM-Series Interfaces
1	Ethernet 1/0 (mgmt)
2	Ethernet 1/1 (eth1)
3	Ethernet 1/2 (eth2)
4	Ethernet 1/3 (eth3)
5	Ethernet 1/4 (eth4)
6	Ethernet 1/5 (eth5)
7	Ethernet 1/6 (eth6)
8	Ethernet 1/7 (eth7)

 ©2024 Palo Alto Networks, Inc.

VMware vNIC	VM-Series Interfaces
9	Ethernet 1/8 (eth8)
10	Ethernet 1/9 (eth9)

The mapping on the VM-Series Firewall remains the same no matter which vNICs you add on ESXi. Interfaces you activate on the firewall always take the next available vNIC on ESXi.

In the following diagram, eth3 and eth4 on the VM-Series Firewall are paired to vNICs 2 and 3 on ESXi, and eth1 and eth2 are unmapped, as shown on the left.

If you want to add two additional interfaces while maintaining the current mapping, activate vNICs 4 and 5 and reboot down the firewall. The existing vNIC mapping is preserved because you added the interfaces after the last-mapped inteface.

If you activate eth1 and eth2 on the VM-Series firewall, the interfaces reorder themselves as shown on the right, resulting in a mapping mismatch that impacts traffic.

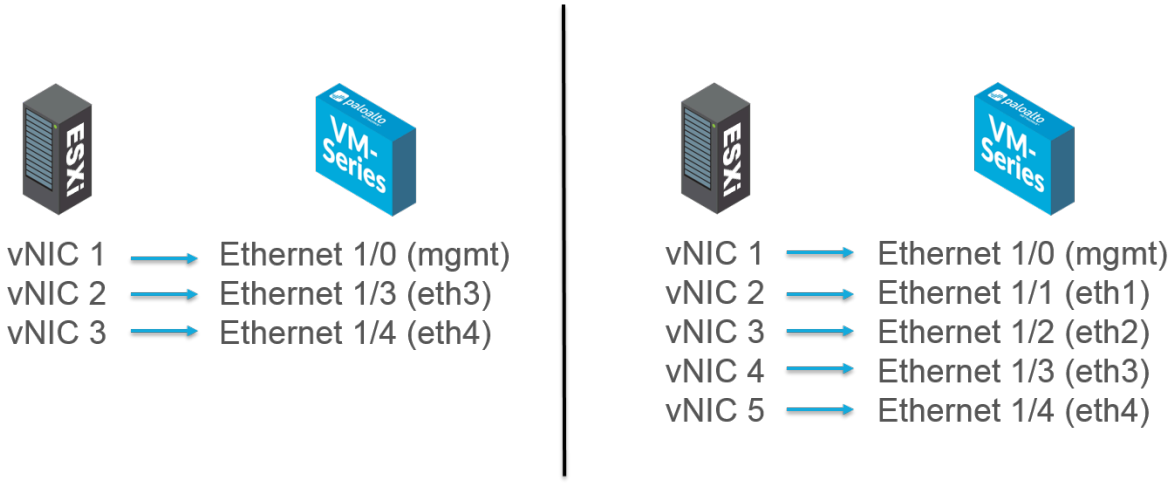

To avoid the issues described in the preceding example, you can do the following:

- When provisioning your ESXi host for the first time, activate all nine vNICs beyond the first. Adding all nine vNICs as placeholders before powering on the VM-Series Firewall allows you to use any VM-Series interfaces regardless of order.
- If all vNICs are active, adding additional interfaces no longer requires a reboot. Because each vNIC on ESXi requires that you choose a network, you can create an empty port group as a network placeholder.
- Do not remove VM-Series firewall vNICs to avoid mapping mismatches.

Provision the VM-Series Firewall on an ESXi Server

Use these instructions to deploy the VM-Series firewall on a (standalone) ESXi server. For deploying the VM-Series NSX edition firewall, see Set Up the VM-Series Firewall on VMware NSX.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Download the OVA file.

Register your VM-Series firewall and obtain the OVA file from the Palo Alto Networks Customer Support web site.

The OVA file contains the base installation. After the base installation is complete, you must download and install the latest PAN-OS version from the support portal. This ensures that you have the latest fixes implemented since the base image was created. For instructions, see Upgrade the PAN-OS Software Version (Standalone Version).

STEP 2 | Before deploying the OVA file, set up virtual standard switch(es) or virtual distributed switch(es) that you need for the VM-Series firewall.

*If you are deploying the VM-Series firewall with Layer 3 interfaces, your firewall uses Hypervisor Assigned MAC Addresses by default. If you choose to disable hypervisor assigned MAC address, or if you are deploying the firewall with Layer 2, virtual wire, or tap interfaces, you must configure (set to **Accept**) any virtual switch attached to the VM-Series firewall to allow the following modes: promiscuous mode, MAC address changes, and Forged transmits.*

Configure a virtual standard switch or a virtual distributed switch to receive frames for the VM-Series firewall.

Virtual Standard Switch

1. Navigate to **Home** > **Hosts and Clusters** and select a host.
2. Click the **Configure** tab and view **Virtual Switches**. For each VM-Series firewall attached a virtual switch, click on **Properties**.
3. Highlight a port group corresponding to a virtual switch and click **Edit Settings**. In the vSwitch properties, click the **Security** tab and set **Promiscuous Mode, MAC Address Changes** and **Forged Transmits** to **Accept** and then click **OK**. This change propagates to all port groups on the virtual switch.

Virtual Distributed Switch

1. Select **Home** > **Networking**. Select your virtual distributed switch and highlight the **Distributed Port Group** you want to edit.
2. Click **Edit Settings**, select **Policies** > **Security**, and set **Promiscuous Mode, MAC Address Changes** and **Forged Transmits** to **Accept** and click **OK**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Deploy the OVA.

If you add additional interfaces (vNICs) to the VM-Series firewall, you must reboot (because new interfaces are detected during the boot cycle). To minimize the need to reboot the firewall, activate the interfaces at initial deployment or during a maintenance window.

*To view the progress of the installation, monitor the **Recent Tasks** list.*

1. Log in to vCenter using the vSphere client. You can also go directly to the target ESXi host if needed.
2. From the vSphere web client, go to **Hosts and Clusters**, right-click your host, and select **Deploy OVF Template**.
3. Browse to the OVA file that you downloaded in 1 Select the file, and click **Next**. Review the template's details and click **Next**.
4. Name the VM-Series firewall instance, and in the **Inventory Location** window, select a Data Center and Folder, and click **Next**.
5. Select an ESXi host for the VM-Series firewall, and click **Next**.
6. Select the datastore to use for the VM-Series firewall, and click **Next**.
7. Leave the default settings for the datastore provisioning, and click **Next**. The default is **Thick Provision Lazy Zeroed**.

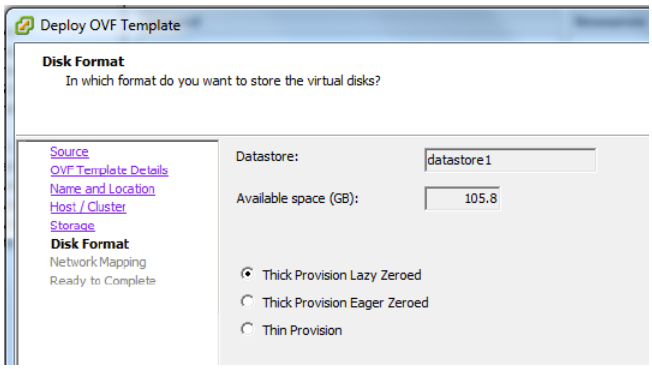

8. Select the networks to use for the two initial vNICs. The first vNIC is used for the management interface and the second vNIC for the first data port. Make sure that the **Source Networks** map to the correct **Destination Networks**.

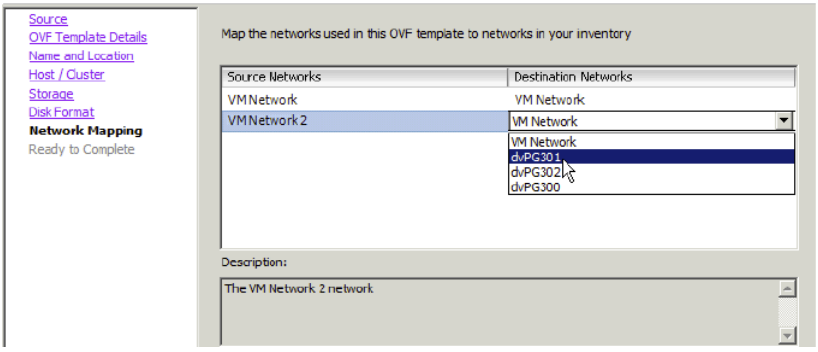

9. Review the details, select **Power on after deployment**, and click **Next**.

 ©2024 Palo Alto Networks, Inc.

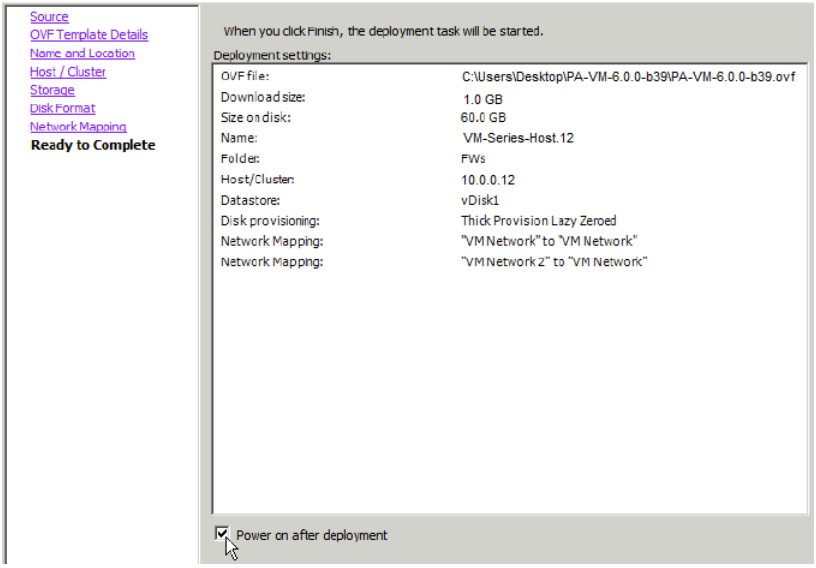

10. When the deployment is complete, click the **Summary** tab to review the current status.

Perform Initial Configuration on the VM-Series on ESXi

Use the virtual appliance console on the ESXi server to set up network access to the VM-Series firewall. By default, the VM-Series firewall uses DHCP to obtain an IP address for the management interface, but, you can also assign a static IP address. After completing the initial configuration, access the web interface to complete further configuration tasks. If you have Panorama for central management, refer to the Panorama Administrator's Guide for information on managing the device using Panorama.

If you are using bootstrapping to perform the configuration of your VM-Series firewall on ESXi, refer to Bootstrap the VM-Series Firewall on ESXi.

For general information about bootstrapping, see Bootstrap the VM-Series Firewall.

STEP 1 | Gather the required information from your network administrator.

- IP address for MGT port
- Netmask
- Default gateway
- DNS server IP address

STEP 2 | Access the console of the VM-Series firewall.

1. Select the **Console** tab on the ESXi server for the VM-Series firewall, or right click the VM-Series firewall and select **Open Console**.
2. Press Enter to access the login screen.
3. Enter the default username/password (admin/admin) to log in.
4. Enter **`configure`** to switch to configuration mode.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure the network access settings for the management interface.

Enter the following commands:

```
set deviceconfig system type static
```

```
set deviceconfig system ip-address <Firewall-IP> netmask <netmask>
default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>
```

STEP 4 | Commit your changes and exit the configuration mode.

Enter **commit**.

Enter **exit**.

STEP 5 | Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server.

1. Use the ping utility to verify network connectivity to the Palo Alto Networks Update server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update server (the Update server does not respond to ping requests.) After verifying DNS resolution, press Ctrl+C to stop the ping request.

```
admin@PA-220 > ping host updates.paloaltonetworks.com
```

```
PING updates.paloaltonetworks.com (10.101.16.13) 56(84) bytes
 of data.
From 192.168.1.1 icmp_seq=1 Destination Host Unreachable
From 192.168.1.1 icmp_seq=2 Destination Host Unreachable
From 192.168.1.1 icmp_seq=3 Destination Host Unreachable
From 192.168.1.1 icmp_seq=4 Destination Host Unreachable
```

2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server: request support check If you have connectivity, the update server responds with the support status for your firewall.

STEP 6 | Apply the capacity auth code and retrieve a license before you begin testing the VM-Series firewall.

An unlicensed VM-Series firewall can process up to approximately 1230 concurrent sessions. Depending on the environment, the session limit can be reached very quickly, causing unpredictable results.

Add Additional Disk Space to the VM-Series Firewall

The VM-Series firewall requires a 60GB virtual disk, of which 21GB is used for logging, by default.

- For large deployments, use Panorama to aggregate data from all next-generation firewalls, and provide visibility across all the traffic on your network. Panorama provides centralized logging and reporting.

 ©2024 Palo Alto Networks, Inc.

- In smaller deployments where you do not use Panorama, you can add a new virtual disk to increase log storage capacity. The new virtual disk can support 60GB to 2TB of storage capacity for logs. This task is described below.

  When the virtual appliance is configured to use a virtual disk, the VM-Series firewall no longer stores logs. If the appliance loses connectivity to the virtual disk, logs can be lost during the failure interval. If necessary, place the newly created virtual disk on a datastore that provides RAID redundancy. RAID10 provides the best write performance for applications with high logging characteristics.

STEP 1 | Power off the VM-Series firewall.

STEP 2 | On the ESXi server, add the virtual disk to the firewall.

1. Select the VM-Series firewall on the ESXi server.
2. Click **Edit Settings**.
3. Click **Add** to launch the Add Hardware wizard, and select the following options when prompted:
   1. Select **Hard Disk** for the hardware type.
   2. Select **Create a new virtual disk**.
   3. Select **SCSI** as the virtual disk type.
   4. Select the **Thick provisioning** disk format.
   5. In the location field, select **Store with the virtual machine option**. The datastore does not have to reside on the ESXi server.
   6. Verify that the settings look correct and click **Finish** to exit the wizard. The new disk is added to the list of devices for the virtual appliance.

STEP 3 | Power on the firewall.

Powering on the firewall initializes the virtual disk for first-time use. The time that the initialization process takes to complete varies by the size of the new virtual disk.

When the new virtual disk is initialized and ready, PAN-OS moves all logs from the existing disk to the new virtual disk. New log entries are now written to this new virtual disk.

PAN-OS also generates a system log entry that records the new disk.

If you reuse a virtual disk that was previously used for storing PAN-OS logs, all logs from the existing disk are overwritten.

STEP 4 | Verify the size of the new virtual disk.

1. Select **Device** > **Setup** > **Management**.
2. In the Logging and Reporting Settings section, verify that the **Log Storage** capacity accurately displays the new disk capacity.

 ©2024 Palo Alto Networks, Inc.

Use VMware Tools on the VM-Series Firewall on ESXi and vCloud Air

The VMware Tools utility improves VM-Series firewall management from vCenter server and vCloud Director. VMware Tools are bundled with the software image for the VM-Series firewall, and all updates are made available with a new OVF image. You cannot manually install or upgrade VMware Tools using the vCenter server or vCloud Director.

- View the IP address(es) on the management interface and the software version on the firewall and Panorama.

  In the Hosts and Cluster section on the vCenter server, select the firewall or Panorama and view the **Summary** tab for information on the IP address(es) assigned to the management interface and the software version currently installed.

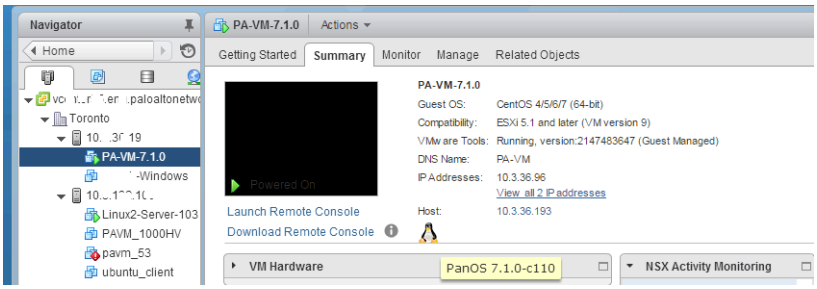

- View resource utilization metrics on hard disk, memory, and CPU. Use these metrics to enable alarms on the vCenter server.

  In the Hosts and Cluster section on the vCenter server, select the firewall or Panorama and view the **Monitor** > **Utilization** tab for information on hard disk, memory, and CPU usage.

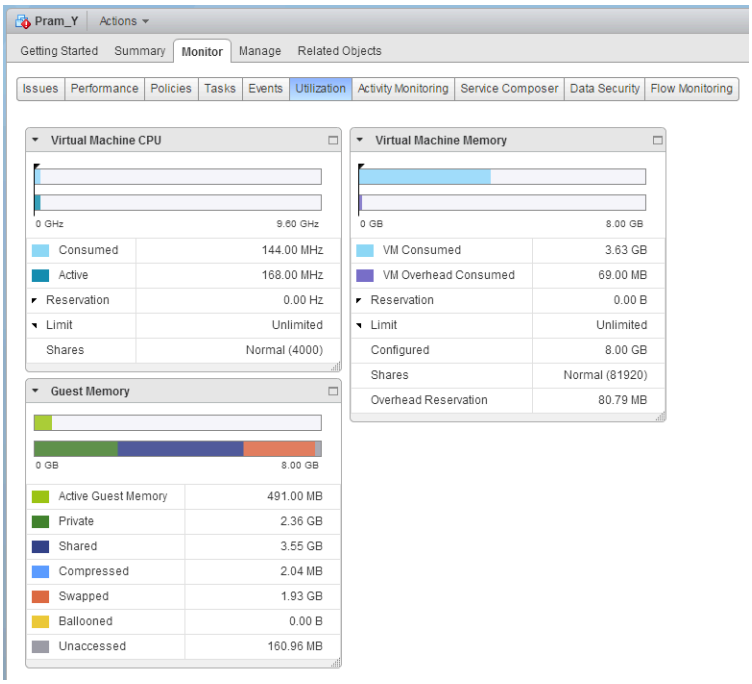

 ©2024 Palo Alto Networks, Inc.

- Gracefully shutdown or restart the firewall and Panorama from the vCenter server.

  In the Hosts and Cluster section on the vCenter server, select the firewall or Panorama and select the **Actions** > **Power** drop-down.

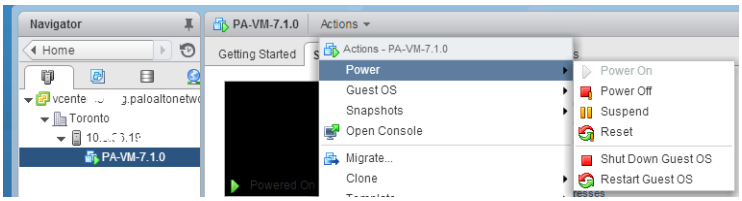

- Create alarm definitions for events you want to be notified about, or events for which you want to specify an automated action.

  Refer to the VMware documentation for details on creating alarm definitions.

  In the Hosts and Cluster section on the vCenter server, select the firewall or Panorama and select the **Manage** > **Alarm Definitions** to add a new trigger and specify an action when a threshold is met. For example, missing heartbeats for a specified duration, or when memory resource usage exceeds a threshold. The following screenshot shows you how to use notifications for heartbeat monitoring on the firewall or Panorama.

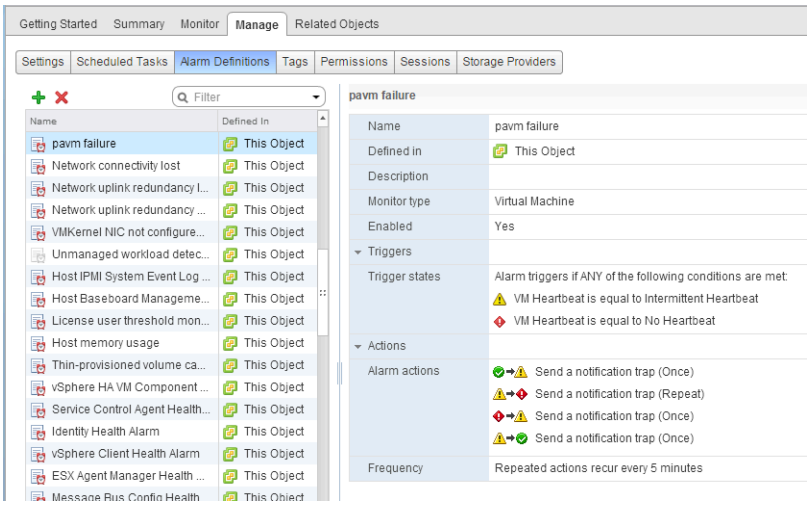

Use vMotion to Move the VM-Series Firewall Between Hosts

To maintain traffic flow while using vMotion to move your VM-Series firewall on VMware ESXi between ESXi hosts with homogeneous CPU configurations, you must use the PAN-OS CLI to pause the internal heartbeat monitoring of the VM-Series firewall during vMotion. You can specify the amount of time, in minutes, that heartbeat monitoring is paused. Heartbeat monitoring can be paused for up to 60 minutes. When the pause interval expires or you deliberately end the pause interval, heartbeat monitoring resumes.

vMotion of the VM-Series firewall is supported on vSphere 6.5, 6.7, and 7.0 if the ESXi hosts have homogeneous CPU configuration. PAN-OS 10.0, 9.1.6, or 9.0.11 and later is required to Use vMotion to Move the VM-Series Firewall Between Hosts installed on vSphere 6.5 or 6.7.

 ©2024 Palo Alto Networks, Inc.

These commands are not required when using vMotion if you are running vSphere 7.0 or later.

STEP 1 | Log in the VM-Series firewall CLI.

STEP 2 | Set the heartbeat monitoring pause interval using the following command. The pause begins as soon as the command is executed. If vMotion is taking longer than expected, you can rerun this command to set a new, longer interval that starts when the command is executed again.

```
request system heartbeat-pause set interval <pause-time-in-minutes>
```

You can view the time remaining in pause interval using the following command.

```
request system heartbeat-pause show interval
```

STEP 3 | (Optional) If you complete vMotion before the pause interval has elapsed, you can end the pause by setting the interval to zero (0).

```
request system heartbeat-pause set interval 0
```

Use the VM-Series CLI to Swap the Management Interface on ESXi

By default, the VM-Series firewall assigns the first interface (eth0) as the management interface. However, in some deployments, the first interface must be pre-mapped to a public IP address. Therefore, the management interface must be assigned to a different interface. Assigning a public IP address to the management interface is a security risk.

This procedure requires VM-Series plugin 2.0.7 or later.

Alternatively, you can enable management interface swap as part of the init-cfg.txt File Components *when bootstrapping.*

STEP 1 | Log in to the VM-Series firewall CLI and enter the following command:

```
set system setting mgmt-interface-swap enable yes
```

STEP 2 | Confirm that you want to swap the interface and use the eth1 dataplane interface as the management interface.

STEP 3 | Reboot the firewall for the swap to take effect. Use the following command:

```
request restart system
```

STEP 4 | Verify that the interfaces have been swapped. Use the following command:

```
debug show vm-series interfaces all
Phoenix_interface   Base-OS_port   Base-OS_MAC        PCI-ID
       Driver
mgt(interface-swap) eth0   0e:53:96:91:ef:29   0000:00:04.0
  ixgbevf
Ethernet1/1           eth1    0e:4d:84:5f:7f:4d       0000:00:03.0
  ixgbevf
```

 ©2024 Palo Alto Networks, Inc.

VM Monitoring on vCenter

Install and configure the Panorama plugin for VMware vCenter to retrieve the IP addresses for guests in your vCenter environment and use that information to build policy using Dynamic Address Groups.

The Panorama plugin for VMware vCenter does not support proxy servers.

- About VM Monitoring on VMware vCenter
- Install the Panorama Plugin for VMware vCenter
- Configure the Panorama Plugin for VMware vCenter

About VM Monitoring on VMware vCenter

The Panorama plugin for VMware vCenter gives you the tools to build policy for your vCenter environment using Dynamic Address Groups. Dynamic address groups allow you to create policy that automatically adapts to changes in your environment, such as the addition or deletion of guests. The VMware vCenter plugin monitors for changes in your vCenter environment and shares that information with Panorama.

The plugin processes the information it receives from vCenter and converts it into a set of tags on Panorama that you can use as match criteria for assigning IP address to dynamic address groups. Each tag has a prefix that describes the hierarchy above the VM.

In this example, each tag in Panorama begins with the prefix shown below. Each tag includes the vCenter name, data center name, and cluster name; if you have folders in your vCenter hierarchy, tags will include the folder names. The order of the objects in the tag matches the order in the vCenter hierarchy.

vcenter.<vcenter-name>_ParentA_ParentB_Datacenter_CHILD1_CHILD2_Cluster_<tag>

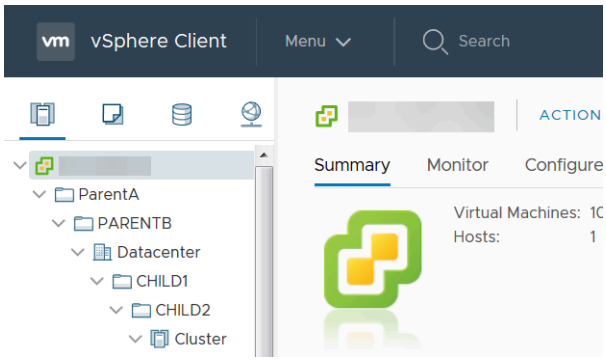

The Panorama plugin for VMware vCenter does not support tags associated to vApps or resource pools.

The tags are shown in Panorama in the following formats:

- **vcenter.<vcenter-name>_<datacenter-name>_<cluster-name>_vmname.<vm-name>**—this tag maps virtual machine IP addresses based on VM name.

 ©2024 Palo Alto Networks, Inc.

- **vcenter.<vcenter-name>_<datacenter-name>_<cluster-name>_guestos.<guest-os>**—this tag maps virtual machine IP addresses based on guest operating system.
- **vcenter.<vcenter-name>_<datacenter-name>_<cluster-name>_annotation.<annotation>**—this tag maps virtual machine IP addresses based on annotation.
- **vcenter.<vcenter-name>_<datacenter-name>_<cluster-name>_vlanId.<vlan-ID>**—this tag maps virtual machine IP addresses based on VLAN ID.
- **vcenter.<vcenter-name>_<datacenter-name>_<cluster-name>_host-ip.<host-ip>**—this tag maps virtual machine IP addresses based on host IP address.
- **vcenter.<vcenter-name>_<datacenter-name>_<cluster-name>_<tag-category>.<user-defined-tag>**—this tag maps virtual machine IP addresses based on user-defined tags created in vCenter.

  The plugin supports a maximum of 16 user-defined tags per VM. Any user-defined tags beyond 16 are not processed.

The Panorama plugin for vCenter cannot process tags that are longer than 128 characters; this includes letters, numbers, and special characters. Whitespace in vCenter object names is replaced with forward slashes. Additionally, Panorama does not support non-ASCII special characters or the following special characters—**'<>&"** in vCenter VM names and annotations. Panorama drops tags containing unsupported characters.

To retrieve endpoint IP-address-to-tag mapping information, you must configure a Monitoring Definition for each vCenter in your virtual environment. The Monitoring Definition specifies the username and password that allows Panorama to connect to vCenter. It also specifies the device groups and corresponding notify groups containing the firewalls to which Panorama pushes the tags. After you configure the Monitoring Definition and the Panorama plugin for VMware vCenter retrieves the tags, you can create DAGs and add the tags as match criteria.

Install the Panorama Plugin for VMware vCenter

To get started with endpoint monitoring on vCenter, download and install the Panorama Plugin for VMware vCenter.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panoramas in an HA pair, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away

 ©2024 Palo Alto Networks, Inc.

- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  ```
  request plugins dau plugin-name <plugin-name> unblock-device-push yes
  ```

  You can cancel this command by executing:

  ```
  request plugins dau plugin-name <plugin-name> unblock-device-push no
  ```

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Select **Panorama** > **Plugins**.

STEP 2 | Select **Upload** and click **Browse** to locate the plugin file.

STEP 3 | Click **OK** to complete the upload.

STEP 4 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

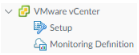

Configure the Panorama Plugin for VMware vCenter

After installing the plugin, complete the following procedure to establish a connection between Panorama and vCenter.

For the plugin to monitor virtual machines in your vCenter environment, you must have VMware tools installed. In vCenter, IP addresses of VMs are not externally retrievable; they are only visible through VMware tools. Additionally, native read-only permissions are required for the plugin to retrieve IP address information from vCenter.

STEP 1 | Log in to the Panorama web interface.

STEP 2 | Enable monitoring and set the monitoring interval.

1. Select **Panorama** > **VMware vCenter** > **Setup** > **General**.
2. Select **Enable Monitoring**. This enables monitoring for all vCenters in your deployment.
3. Set the **Monitoring Interval** in seconds. The monitoring interval is how often Panorama retrieves updated network information from vCenter. The default value is 60 seconds and has a range of 60 to 84600 seconds.

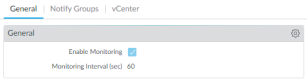

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create a notify group.

1. Select **Panorama** > **VMware vCenter** > **Setup** > **Notify Groups**.
2. Click **Add**.
3. Enter a descriptive **Name** for your notify group.
4. Select the device groups in your vCenter deployment.

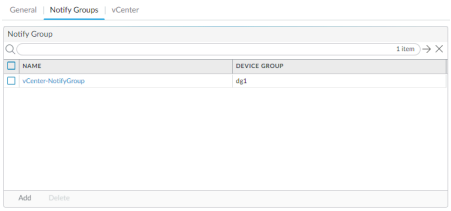

STEP 4 | Add vCenter information. The Panorama plugin for VMware vCenter supports up to 16 vCenter instances.

1. Select **Panorama** > **VMware vCenter** > **Setup** > **vCenter**.
2. Enter a descriptive **Name** for your vCenter.
3. Enter the IP address or FQDN for vCenter and port, if applicable.
4. Enter your vCenter username.
5. Enter and confirm your vCenter password.
6. Click **Validate** to verify that Panorama can connect to vCenter using the login credentials you entered.
7. Click **OK**.

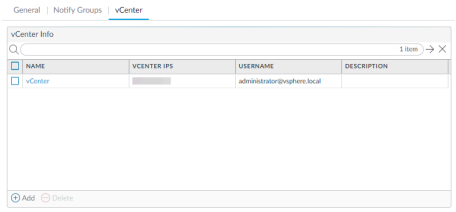

STEP 5 | Configure up to 16 Monitoring Definitions.

A vCenter instance can be assigned to only one Monitoring Definition.

1. Select **Panorama** > **VMware vCenter** > **Monitoring Definition** and click **Add**.
2. Enter a descriptive **Name** and optionally a description to identify the vCenter for which you use this definition.
3. Select the **vCenter** and **Notify Group**.
4. Click **OK**.

STEP 6 | **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Verify that you can view the VM information on Panorama, and define the match criteria for Dynamic Address Groups.

You must use the OR operator when using more than one tag in the match criteria; using the AND operator does not work.

Some browser extensions may block API calls between Panorama and vCenter which prevents Panorama from receiving match criteria. If Panorama displays no match criteria and you are using browser extensions, disable the extensions and Synchronize Dynamic Objects to populate the tags available to Panorama.

STEP 8 | Verify that addresses in your VMs are added to DAGs.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click **More** in the Addresses column of a DAG.

   Panorama displays a list of IP addresses added to that DAG based on the match criteria you specified.

STEP 9 | Use dynamic address groups in policy.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

STEP 10 | You can update the dynamic objects from vCenter at any time by synchronizing dynamic objects. Synchronizing dynamic objects enables you to maintain context on changes in the virtual environment and allows you to enable applications by automatically updating the Dynamic Address Groups used in policy rules.

1. Select **Panorama** > **VMware vCenter** > **Monitoring Definition**.
2. Click **Synchronize Dynamic Objects**.

STEP 11 | If a firewall in your vCenter deployment restarts or disconnects from Panorama, that firewall goes out of sync with the Panorama plugin for vCenter and no receive updates. After the firewall reconnects with Panorama, you must manually synchronize Panorama and the firewall.

1. Log in to the Panorama CLI.
2. Execute the following command.

   ```
   admin@Panorama> request plugins vmware_vcenter sync
   ```

 ©2024 Palo Alto Networks, Inc.

Troubleshoot ESXi Deployments

Many of the troubleshooting steps for the VM-Series firewall are very similar to the hardware versions of PAN-OS. When problems occur, you should check interface counters, system log files, and if necessary, use debug to create captures.

The following sections describe how to troubleshoot some common problems:

- Basic Troubleshooting
- Installation Issues
- Licensing Issues
- Connectivity Issues

Basic Troubleshooting

Recommendation for Network Troubleshooting Tools

It is useful to have a separate troubleshooting station to capture traffic or inject test packets in the virtualized environment. It can be helpful to build a fresh OS from scratch with common troubleshooting tools installed such as tcpdump, nmap, hping, traceroute, iperf, tcpedit, netcat, etc. This machine can then be powered down and converted to a template. Each time the tools are needed, the troubleshooting client (virtual machine) can be quickly deployed to the virtual switch(es) in question and used to isolate networking problems. When the testing is complete, the instance can simply be discarded and the template used again the next time it is required.

For performance related issues on the firewall, first check the **Dashboard** from the firewall web interface. To view alerts or create a tech support or stats dump files navigate to **Device** > **Support**.

For information in the vSphere client go to **Home** > **Inventory** > **VMs and Templates**, select the VM-Series firewall instance and click the **Summary** tab. Under **Resources**, check the statistics for consumed memory, CPU and storage. For resource history, click the **Performance** tab and monitor resource consumption over time.

Installation Issues

- Issues with Deploying the OVA
- Why does the firewall boot into maintenance mode?
- How do I modify the base image file for the VM-1000-HV license?

Issues with Deploying the OVA

- The VM-Series is delivered as a zip archive in the Open Virtualization Alliance (OVA) format that expands into three files.

  If you are having trouble deploying the OVA image, make sure the three files are unpacked and accessible. If necessary, download and extract the OVA image again.

 ©2024 Palo Alto Networks, Inc.

- The virtual disk in the OVA image is nearly 1GB. It must be present on the computer running the vSphere client, or it must be accessible as a URL for the OVA image.
  - Make sure the network connection between the vSphere client computer and the target ESXi host has low latency and sufficient bandwidth. If the connection is poor, the OVA deployment can take hours, or timeout and fail.

    You can minimize this problem if you host the image on a device in the same network as the ESXi host.
  - Any firewalls in the path must allow TCP ports 902 and 443 from the vSphere client to the ESXi host(s).
- ESX 6.5.0a build 4887370 limits you to 2 CPU cores per socket. If you are deploying a VM-300, VM-500 or VM-700 to which you want to allocate more than 2 vCPUs per socket, refer to the VMware KB: https://kb.vmware.com/s/article/53354, for a workaround.

Why does the firewall boot into maintenance mode?

If you have purchased the VM-1000-HV license and are deploying the VM-Series firewall in standalone mode on a VMware ESXi server, you must allocate the minimum memory your VM-Series model requires.

To avoid booting in maintenance mode, you must either modify the base image file (see How do I modify the base image file for the VM-1000-HV license?), or, edit the settings on the ESXi host or the vCenter server before you power on the VM-Series firewall.

Also, verify that the interface is VMXnet3. Setting the interface type to any other format causes the firewall to boot into maintenance mode.

How do I modify the base image file for the VM-1000-HV license?

If you have purchased the VM-1000-HV license and are deploying the VM-Series firewall in standalone mode on a VMware ESXi server, use these instructions to modify the following attributes that are defined in the base image file (.ova or .xva) of the VM-Series firewall.

Important: Modifying values other than those listed here invalidates the base image file.

STEP 1 | Open the base image file, for example 7.0.0, with a text editing tool such as notepad.

STEP 2 | Search for 4096 and change the memory allocated to 5012 (that is 5 GB) as follows:

```
<Item>
        <rasd:AllocationUnits>byte * 2^20</rasd:AllocationUnits>
        <rasd:Description>Memory Size</rasd:Description>
        <rasd:ElementName>4096MB of memory</rasd:ElementName>
        <rasd:InstanceID>2</rasd:InstanceID>
        <rasd:ResourceType>4</rasd:ResourceType>
        <rasd:VirtualQuantity>4096</rasd:VirtualQuantity>
<Item>
        <rasd:AllocationUnits>byte * 2^20</rasd:AllocationUnits>
        <rasd:Description>Memory Size</rasd:Description>
        <rasd:ElementName>5120MB of memory</rasd:ElementName>
        <rasd:InstanceID>2</rasd:InstanceID>
        <rasd:ResourceType>5</rasd:ResourceType>
        <rasd:VirtualQuantity>5120</rasd:VirtualQuantity>
```

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Change the number of virtual CPU cores allotted from 2 to 4 or 8 as desired for your deployment:

```
<Item>
        <rasd:AllocationUnits>hertz * 10^6</rasd:AllocationUnits>
        <rasd:Description>Number of Virtual CPUs</
rasd:Description>
        <rasd:ElementName>2 virtual CPU(s)</rasd:ElementName>
        <rasd:InstanceID>1</rasd:InstanceID>
        <rasd:ResourceType>3</rasd:ResourceType>
        <rasd:VirtualQuantity>2</rasd:VirtualQuantity>
        <vmw:CoresPerSocket ova:required="false">2</
vmw:CoresPerSocket>
      </Item>
<Item>
        <rasd:AllocationUnits>hertz * 10^6</rasd:AllocationUnits>
        <rasd:Description>Number of Virtual CPUs</
rasd:Description>
        <rasd:ElementName>4 virtual CPU(s)</rasd:ElementName>
        <rasd:InstanceID>1</rasd:InstanceID>
        <rasd:ResourceType>3</rasd:ResourceType>
        <rasd:VirtualQuantity>4</rasd:VirtualQuantity>
        <vmw:CoresPerSocket ova:required="false">2</
vmw:CoresPerSocket>
      </Item>
```

Alternatively, you can deploy the firewall, and before you power on the VM-Series firewall, edit the memory and virtual CPU allocation directly on the ESXi host or the vCenter server.

Licensing Issues

- Why am I unable to apply the support or feature license?
- Why does my cloned VM-Series firewall not have a valid license?
- Does moving the VM-Series firewall cause license invalidation?

Why am I unable to apply the support or feature license?

Have you applied the capacity auth-code on the VM-Series firewall? Before you can activate the support or feature license, you must apply the capacity auth-code so that the device can obtain a serial number. This serial number is required to activate the other licenses on the VM-Series firewall.

Why does my cloned VM-Series firewall not have a valid license?

VMware assigns a unique UUID to each virtual machine including the VM-Series firewall.So, when a VM-Series firewall is cloned, a new UUID is assigned to it. Because the serial number and license for each instance of the VM-Series firewall is tied to the UUID, cloning a licensed VM-Series firewallresults in a new firewall with an invalid license. You need a new auth-code to activate the license on the newly deployed firewall. You must apply the capacity auth-code and a new support license in order to obtain full functionality, support, and software upgrades on the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

Does moving the VM-Series firewall cause license invalidation?

If you are manually moving the VM-Series firewall from one host to another, be sure to select the option, **This guest was moved** to prevent license invalidation.

Connectivity Issues

- Why is the VM-Series firewall not receiving any network traffic?

Why is the VM-Series firewall not receiving any network traffic?

On the VM-Series firewall. check the traffic logs (**Monitor** > **Logs**). If the logs are empty, use the following CLI command to view the packets on the interfaces of the VM-Series firewall:

```
show counter global filter
delta yes
Global counters:
Elapsed time since last sampling: 594.544 seconds
---------------------------------------------------------------------------
Total counters shown: 0
---------------------------------------------------------------------------
```

In the vSphere environment, check for the following issues:

- Check the port groups and confirm that the firewall and the virtual machine(s) are on the correct port group

  Make sure that the interfaces are mapped correctly.

  Network adapter 1 = management

  Network adapter 2= Ethernet1/1

  Network adapter 3 = Ethernet1/2

  For each virtual machine, check the settings to verify the interface is mapped to the correct port group.

 ©2024 Palo Alto Networks, Inc.

- Verify that either promiscuous mode is enabled for each port group or for the entire switch or that you have configured the firewall to Hypervisor Assigned MAC Addresses.

  Since the dataplane PAN-OS MAC addresses are different than the vNIC MAC addresses assigned by vSphere, the port group (or the entire vSwitch) must be in promiscuous mode if not enabled to use the hypervisor assigned MAC address:

  - Check the VLAN settings on vSphere.

    The use of the VLAN setting for the vSphere port group serves two purposes: It determines which port groups share a layer 2 domain, and it determines whether the uplink ports are tagged (802.1Q).

  - Check the physical switch port settings

    If a VLAN ID is specified on a port group with uplink ports, then vSphere uses 802.1Q to tag outbound frames. The tag must match the configuration on the physical switch or the traffic does not pass.

    Check the port statistics if using virtual distributed switches (vDS); Standard switches do not provide any port statistics

 ©2024 Palo Alto Networks, Inc.

Performance Tuning of the VM-Series for ESXi

The VM-Series firewall for ESXi is a high-performance appliance but may require tuning of the hypervisor to achieve the best results. This section describes some best practices and recommendations for facilitating the best performance of the VM-Series firewall. For the best performance, ESXi 6.0.0.0 or later is recommended.

- Install the NIC Driver on ESXi
- Enable DPDK on ESXi
- Enable SR-IOV on ESXi
- Enable ESXi VLAN Access Mode with SR-IOV
- Enable Multi-Queue Support for NICs on ESXi
- VNF Tuning for Performance

Install the NIC Driver on ESXi

For the best performance, use SR-IOV with Intel 10GB network interfaces which requires the ixgbe 4.4.1 driver to support multiple queues for each interface.

STEP 1 | Obtain a list of network interfaces on the ESXi host.

1. Log in to the ESXi host CLI.
2. Use the following command to return a list of network interfaces:

```
$ esxcli network nic list
```

STEP 2 | Determine the driver version for a particular interface.

You can use either **ethtool** or **esxcli** to determine the currently-installed driver version. The following example uses vNIC4 and returns driver version 3.21.6.

- ethtool—**ethtool -l *<nic-name>***

```
$ ethtool -I vNIC4
driver: ixgbe
version: 3.21.6iov
firmware-version: 0x80000389
bus-info: 0000:04:00.0
```

- esxcli—**esxcli network nic get -n *<nic-name>***

```
$  esxcli network nic get -n vNIC4
   Advertised Auto Negotiation: true
   Advertised Link Modes:
   Auto Negotiation: true
   Cable Type:
   Current Message Level: 7
   Driver Info:
         Bus Info: 0000:04:00.0
```

 ©2024 Palo Alto Networks, Inc.

```
        Driver: ixgbe
        Firmware Version: 0x80000389
        Version: 3.21.6iov
   Link Detected: false
   Link Status: Down
   Name: vNIC4
   PHYAddress: 0
   Pause Autonegotiate: true
   Pause RX: true
   Pause TX: true
   Supported Ports: FIBRE
   Supports Auto Negotiation: true
   Supports Pause: true
   Supports Wakeon: false
   Transceiver: external
   Wakeon: None
```

STEP 3 | Install the new driver.

1. Download the ixgbe 4.4.1 driver from the VMware website. Extract the contents to a local directory and find the .zip or .vib files for your driver.
2. Create a new folder in your ESXi host datastore.
3. Copy the local .zip or .vib file you extracted to the new folder in your ESXi host datastore.
4. Enable maintenance mode on the ESXi host.
5. Use one of the following commands to install the new driver, using -d for .zip files, or -v for .vib files.

   - `$ esxcli software vib install -d <path to driver .zip file>`
   - `$ esxcli software vib install -v <path to driver .vib file>`

   You must specify the absolute path to the .zip or .vib file. For example:

   ```
   $ esxcli software vib install -d  "/vmfs/volumes/
   Datastore/DirectoryName/DriverName.zip"
   ```

6. Verify the VIB installation.

   ```
   $ esxcli software vib list
   ```

7. Reboot the ESXi host.

Enable DPDK on ESXi

The Data Plane Development Kit (DPDK) enhances VM-Series performance by increasing network interface card (NIC) packet processing speed. On the VM-Series firewall, DPDK is enabled by default on ESXi.

To support DPDK, all data interfaces must be using the same driver.

 ©2024 Palo Alto Networks, Inc.

To take advantage of DPDK, you must use a NIC with one of the supported DPDK drivers mentioned in DPDK Driver Versions:

If you disable DPDK, the NIC uses PacketMMap instead of DPDK. You can disable DPDK using the command **`set system setting dpdk-pkt-io off`**.

See the Compatibility Matrix for ESXi hypervisor support and PacketMMAP and DPDK driver support by PAN-OS version.

Enable SR-IOV on ESXi

Single root I/O virtualization (SR-IOV) allows a single PCIe physical device under a single root port to appear to be multiple separate physical devices to the hypervisor or guest. Enable SR-IOV by enabling virtual function devices on the SR-IOV NIC and the modify the guest settings in vCenter.

SR-IOV on the VM-Series for ESXi requires one of the Intel NIC drivers mentioned in PacketMMAP Driver Versions. See the Compatibility Matrix for SR-IOV and DPDK driver support by PAN-OS version.

There are two ways to enable SR-IOV on ESXi.

- **SR-IOV passthrough**—In this method you enable virtual function devices on the SR-IOV NIC and modify the guest settings in vCenter, adding the SR-IOV VF interface as adaptor type “SR-IOV passthrough”. Refer to Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine.

  This method, which is preferred for PAN-OS 8.1.2 and later, allows you to add the SR-IOV PF to a vSwitch or DvSwitch.
- **PCI Adaptor**—This method was required for PAN-OS 8.0 through 8.1.1. You can view the PCI Adaptor workflow in Enable SR-IOV on ESXi in the 8.1 Deployment Guide.

  The PCI Adaptor method has the limitation that you cannot configure a vSwitch on the physical port on which you enable SR-IOV. The VM-Series firewall must have exclusive access to the physical port and associated virtual functions (VFs) on that interface so it can communicate with the host or other virtual machines on the network. Refer to Add a PCI Device in the vSphere Web Client.

Enable ESXi VLAN Access Mode with SR-IOV

The VM-Series firewalls on ESXi can operate in VLAN access mode to support use cases where it is deployed as a virtual network function (VNF) that offers security-as-a-service in a multi-tenant cloud/data center environment. In VLAN access mode, each VNF has dedicated virtual network interfaces (VNIs) for each network and it sends and receives packets to/from SR-IOV virtual functions (VFs) without VLAN tags; you must enable this capability on the physical and virtual functions on the host hypervisor. When you, then enable VLAN access mode on the VM-Series firewall, the firewall can send and receive traffic without VLAN tags across all its dataplane interfaces. Additionally, if you configure QoS policies, the firewall can enforce QoS on the access interface and provide differentiated treatment of traffic in a multi-tenant deployment.

By default, the VM-Series firewall on ESXi operates in VLAN trunk mode.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | On the host system, set up the physical and virtual function to operate in VLAN access mode.

1. Click **Networking** in the VMware Host Client inventory and click **Port groups**.
2. In the list that you want to edit, right-click the port group and select **Edit settings**.Enter a new port group **Name**.Enter a new value for the **VLAN ID**.

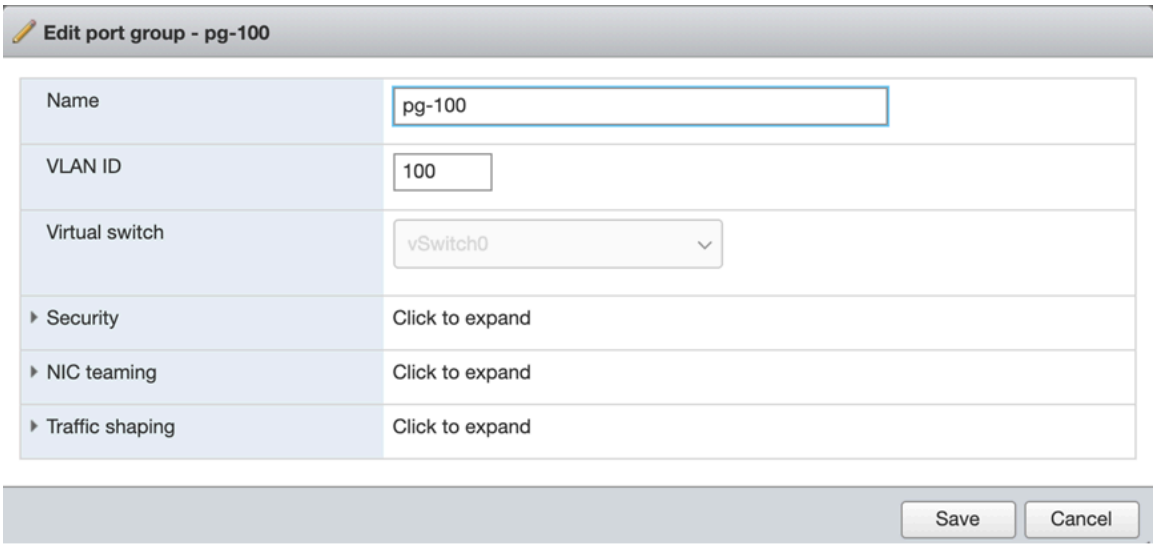

For best performance on the VM-Series firewall, make sure to:

- *Enable CPU pinning.*
- *Disable Replay Protection, if you have configured IPSec Tunnels.*

  *On the firewall web interface, select **Network** > **IPSec Tunnels**, select an IPSec tunnel, click **General**, select **Show Advanced Options**, and clear **Enable Replay Protection**.*

STEP 2 | Access the CLI on the VM-Series firewall.

STEP 3 | Enable VLAN access mode.

```
request plugins vm-series vlan-mode access-mode on
```

on enables VLAN access mode; to use VLAN trunk mode, enter **request plugins vm-series vlan-mode access-mode off**.

STEP 4 | Reboot the firewall.

```
request restart system
```

STEP 5 | Verify the VLAN mode configuration.

```
show plugins vm-series vlan-mode
```

Enable Multi-Queue Support for NICs on ESXi

Multi-queue allows network performance to scale with the number of vCPUs and allows for parallel packet processing by creating multiple TX and RX queues. Modify the .vmx file or access Advanced Settings to enable multi-queue.

 ©2024 Palo Alto Networks, Inc.

The pNIC setting is also applicable for NSX-T since ESXi is the hypervisor for NSX-T deployments.

STEP 1 | Enable multi-queue.

1. Open the .vmx file.
2. Add the following parameter:

```
ethernetX.pnicFeatures = “4”
```

STEP 2 | Enable receive-side scaling (RSS).

1. Log in to the CLI on the ESXi host.
2. Execute the following command:

```
$ vmkload_mod -u ixgbe
$ vmkload_mod ixgbe RSS=”4,4,4,4,4,4”
```

STEP 3 | For the best performance, allocate additional CPU threads per ethernet/vSwitch device. This is limited by the amount of spare CPU resources available on the ESXi host.

1. Open the .vmx file.
2. Add the following parameter:

```
ethernetX.ctxPerDev = “1”
```

VNF Tuning for Performance

This topic provides VNF tuning guidance for VM-Series deployments. It is a reference to help you choose some of the parameter settings for a VM-Series deployment. Before attempting tuning, you should be familiar with the steps to install the VM-Series firewall on the VMware vSphere hypervisor (ESXi), including how to configure tuning parameters and attributes.

This guidance might not apply to VM-Series deployments on top of white-box or grey-box environments targeting SD-WAN, MSSP, or CSSP use-cases.

VM-Series is a high-performance appliance and is available in various form-factors depending on size, hypervisor footprint, and its deployment location in either private or public cloud.

Global and host-level configuration changes impact other VMs running on the same host. You should consider any trade-offs and prudently choose the parameters that best suit your deployment.

- ESXi Tuning Parameters
- Use Cases
- References

 ©2024 Palo Alto Networks, Inc.

ESXi Tuning Parameters

To achieve best results in performance on VM-series, you can tune hardware, hypervisor, and network I/O parameters.

The parameters mentioned here do not apply to every deployment model.

- BIOS Settings
- Physical Settings
- Virtual NIC Settings
- NUMA and Resource Considerations

BIOS Settings

This section recommends BIOS Power Management, Hyperthreading, and Intel VT-D settings that can enhance VM-Series firewall performance, and concludes with a sample BIOS configuration.

- Power Management
- Hyperthreading
- Intel Virtualization Technology for Directed I/O
- Sample BIOS Configuration

Power Management

For latency-sensitive applications, any form of power management adds latency to the path where an idle system (in one of several power-saving modes) responds to an external event. VMware recommends setting the BIOS power management setting to “static high performance” (no OS-controlled power management), effectively disabling any form of active power management. Servers with Intel Nehalem class and later CPUs (Intel Xeon 55xx and later) offer two other power management options: C-states and Intel Turbo Boost.

Leaving C-states enabled can increase memory latency and is therefore not recommended for low-latency workloads. Even the enhanced C-state, known as C1E, introduces longer latencies to wake up the CPUs from halt (idle) states to full-power. VMware recommends disabling C1E in the BIOS to further lower latencies.

- For HP, set Power Regulator Mode to Static High Mode and disable QPI Processor, C-state support, and C1E Support.
- For Dell, set Power Management Mode, CPU power, and Performance Management to Maximum Performance.

Another parameter to consider is P-states. For outright performance considerations, disable P-state settings on BIOS.

Intel Turbo Boost can lead to performance variations over a period of time. For consistent and deterministic performance, disable Turbo Boost.

Hyperthreading

If the hardware and BIOS support hyperthreading, ESXi automatically enables hyperthreading on hosts. For the best performance from VM series firewalls, disable hyperthreading on ESXi hosts.

 ©2024 Palo Alto Networks, Inc.

If the deployment environment warrants enabling hyperthreading, then ensure that all CPU resources for the VM-Series firewall are reserved from the same NUMA/Socket node that has access to the PCI devices.

In general, configure the PA-VM as a single NUMA VM. See NUMA and Resource Considerations for more details.

Intel Virtualization Technology for Directed I/O

Intel Virtualization Technology for Directed I/O (Intel VT-D) allows a LAN card to be dedicated to a guest system, which enables increased network performance beyond that of an emulated LAN card. Enable this feature at the BIOS. If you plan to leverage SR-IOV for performance (recommended), enable the SRI-OV BIOS setting.

Sample BIOS Configuration

The following screenshots show the system profile settings and the processor settings for a Dell BIOS.

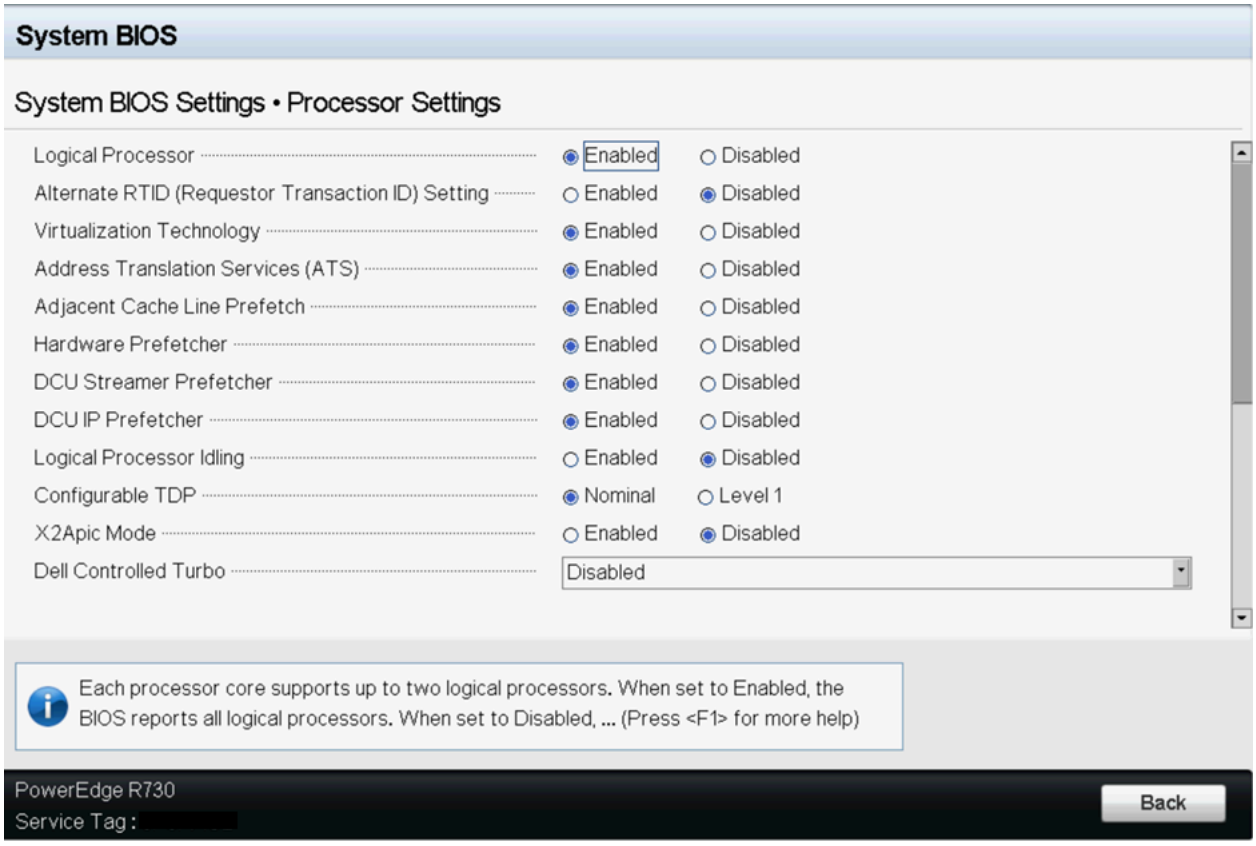

 ©2024 Palo Alto Networks, Inc.

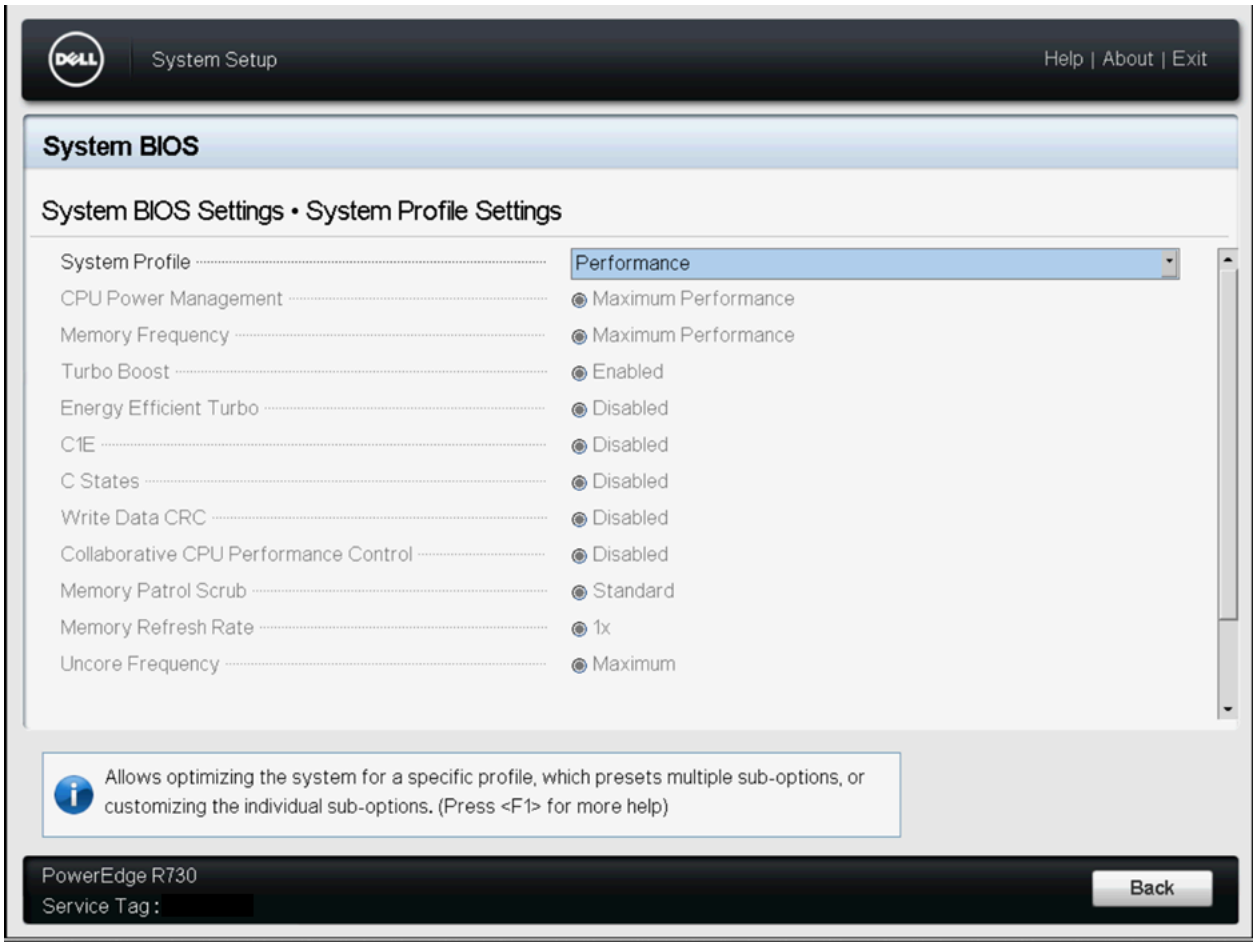

 ©2024 Palo Alto Networks, Inc.

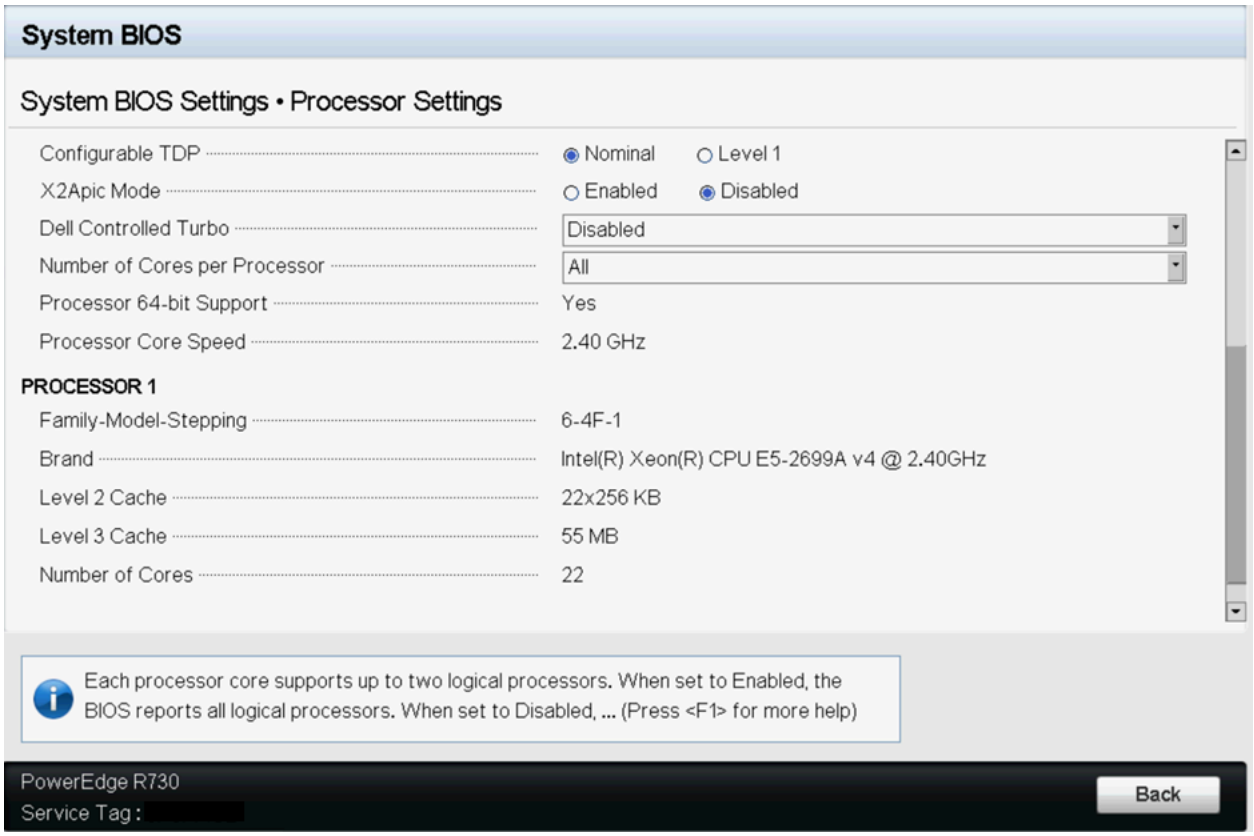

Physical Settings

Most 1GbE or 10GbE network interface cards (NICs) support a feature called interrupt moderation or interrupt throttling, which coalesces interrupts from the NIC to the host so that the host doesn't get overwhelmed and spend all its CPU cycles processing interrupts. However, for latency-sensitive workloads, the time the NIC is delaying the delivery of an interrupt for a received packet or a packet that has successfully been sent on the wire is the time that increases the latency of the workload. For best performance on PA-VM, disable interrupt moderation. For example, disable physical NIC interrupt moderation on the ESXi host as follows:

```
# esxcli system module parameters set -m ixgbe -p
 "InterruptThrottleRate=0"
```

- Transmit Queue
- Queue Pairing

Transmit Queue

The ESXi uplink pNIC layer also maintains a software Tx queue of packets queued for transmission, which by default holds 500 packets. If the workload is I/O intensive with large bursts of transmit packets, this queue can overflow, leading to packets being dropped in the uplink layer. The Tx queue size can be increased up to 10,000 packets with the following ESXi command:

```
# esxcli system settings advanced set -i 10000 -o /Net/
MaxNetifTxQueueLen
```

 ©2024 Palo Alto Networks, Inc.

Depending on the physical NIC and the specific version of the ESXi driver being used on the ESXi host, sometimes packets can be dropped in the pNIC driver because the transmit ring on the pNIC is too small and is filled up. Most pNIC drivers allow you to increase the size of the transmit ring using the following command:

```
# ethtool -G vmnic0 tx 4096
```

This command increases the Tx ring size to 4096 entries. The maximum size you can set for a specific pNIC driver, as well as the current Tx ring size in effect, can be determined using the following command:

```
# ethtool -g vmnic0
```

```
Ring parameters for vmnic0:

Pre-set maximums:
RX: 4096
RX Mini: 0
RX Jumbo: 0
TX: 4096
Current hardware settings:
RX: 512
RX Mini: 0
RX Jumbo: 0
TX: 4096
```

Queue Pairing

Some pNIC drivers, such as Intel's ixgbe and Broadcom's bnx2x, also support "queue pairing", which indicates to the ESXi uplink layer that the receive thread (NetPoll) will also process completion of transmitted packets on a paired transmit queue. For certain transmit-heavy workloads, this can cause delays in processing transmit completions, causing the transmit ring for the vNIC to run out of room for transmitting additional packets, and forcing the vNIC driver in the guest OS to drop packets.

Disabling queue pairing for all pNICs on an ESXi host creates a separate thread for processing pNIC transmit completions. As a result, completions are processed in a timely manner, freeing space in the vNIC's transmit ring to transmit additional packets.

The ESXi command to disable queue pairing is:

```
# esxcli system settings advanced set -o /Net/
NetNetqRxQueueFeatPairEnable
  -i 0
```

For this to take effect, you must reboot the ESXi host.

If PCI-pass through on VM-700 is used on a dedicated host, no performance tuning of the NIC/NIC driver is needed. However, this deployment mode is not common.

 ©2024 Palo Alto Networks, Inc.

Virtual NIC Settings

If possible, use SR-IOV for better performance, as explained in the following topics:

- SR-IOV
- VMXNET3/vSwitch and Virtual Interrupt Coalescing
- Enable Multiqueue Support on Intel x710/x520

SR-IOV

- Changing module parameters for an SR-IOV driver requires an ESXi host reboot.
- Disable physical NIC interrupt moderation on ESXi host as follows:

```
# esxcli system module parameters set -m ixgbe -p
 "InterruptThrottleRate=0“
```

- If you enable multiqueue support, you must also enable Receive-Side Scaling (RSS) for the driver.
  - To enable RSS, set the port value to 4.
  - Specify ports in a comma-separated string.

  Example—Set 3 NICs with 2 ports each.

```
$ vmkload_mod -u ixgbe esxcli system module parameters set -m
 ixgbe
  -p RSS=”4,4,4,4,4,4”
```

```
$ vmkload_mod ixgbe RSS=”4,4,4,4,4,4”
```

  Example—Set RSS for a single port:

```
$ vmkload_mod -u ixgbe esxcli system module parameters set -m
 ixgbe
  -p RSS=”0,4,0,0,0,0”
```

VMXNET3/vSwitch and Virtual Interrupt Coalescing

By default, VMXNET3 supports an interrupt coalescing algorithm (for the same reasons that physical NICs implement interrupt moderation). To avoid flooding the host system with too many interrupts, packets are collected and one single interrupt is generated for multiple packets. This is called interrupt coalescing.

Interrupt coalescence refers to the amount of traffic that a network interface receives, or the amount of time that passes after traffic is received, before you issue a hard interrupt. Interrupting too soon or too frequently results in poor system performance, as the kernel stops (or “interrupts”) a running task to handle the interrupt request from the hardware. Interrupting too late can result in traffic loss if the traffic is not taken off the NIC soon enough—more traffic arrives, overwriting the previous traffic still waiting to be received into the kernel.To disable this functionality through the vSphere Web Client, go to **VM Settings** > **Options** > **Advanced General**

 ©2024 Palo Alto Networks, Inc.

> **Configuration Parameters** and add an entry for **`ethernetX.coalescingScheme`** with the value **`disabled`**.

To disable virtual interrupt coalescing for all virtual NICs on the host (which affects all VMs, not just the latency-sensitive ones), set the advanced networking performance option. Go to **Configuration** > **Advanced Settings** > **Net** and set **`CoalesceDefaultOn`** to **`0`** (disabled).

Enable Multiqueue Support on Intel x710/x520

Use ESXi 6.0.0 or later, with an ixgbe driver version with multiqueue support. See SR-IOV Driver Versions in the Compatibility Matrix. Modify the `.vmx` file or access **Advanced Settings** to enable multiqueue support:

```
ethernetX.pnicFeatures = “4”
```

To set multi-core affinity so a vSwitch can exceed 300K PPS, set:

```
ethernetX.pnicFeatures = "4"
ethernetX.ctxPerDev = "1"
```

Setting **`ethernetX.ctxPerDev = “1”`**, is like a binary flag (set to 1 to enable). This binary flag adds a CPU thread to process traffic only from the port **`ethernetX`**. This leads to improved traffic scheduling performance.

NUMA and Resource Considerations

NUMA is Non-Uniform Memory Access. Multi-Core processors have complicated designs. To tackle performance issues in such systems, you need to be aware of all NUMA and CPU Pinning nuances. Vital aspects to look for:

- Which cores are our threads are running on? (if hyperthreading is enabled, check Hyperthreading)
- Which cores are our vCPUs are running on? (affinity)
- In which NUMA socket is the physical NIC card installed?
- Where has memory been allocated? (NUMA effects)

  Threads running on any socket see one unified memory space – therefore they can read/write to memory that is local to other Sockets.
- Is memory shared between different sockets on a node?
- It takes more time to access memory on different sockets than it takes to access local memory.

  NUMA effects occur when threads excessively access memory on a different NUMA domain. To avoid cross-NUMA issues, avoid Quick Path Interconnect (QPi) between Socket 0 communication and Socket 1.

For latency-sensitive VMs like PA-VM, VMware recommends that you do not over-commit vCPUs as compared to the number of physical CPUs (processors) on the ESXi host. For example, if the host has 8 CPU cores, limit the number of vCPUs for your VM to 7. This ensures that the ESXi VMkernel scheduler has a better chance of placing the vCPUs on pCPUs that won’t contend with other scheduling contexts, such as vCPUs from other VMs or ESXi helper worlds. It is a good

 ©2024 Palo Alto Networks, Inc.

practice to ensure that the number of vCPUs you allocate to the VM does not exceed the number of active CPU-consuming processes or threads in the VM.

For best performance, all vCPUs should be scheduled on the same NUMA node and all VM memory should fit and be allocated out of the local physical memory attached to that NUMA node. This can be changed using the VM setting **`numa.nodeAffinity=0, 1, …`** where 0, 1, and so forth, are the socket numbers.

To ensure that the VM gets exclusive access to the CPU resources, set Latency Sensitivity to High. For the new setting to take effect, the VM CPU reservation must be set to maximum, Memory should be reserved, and the CPU limit must be set to unlimited.

- In newer versions, use the vSphere Web Client to set the VM Latency Sensitivity option to High (the default is Normal).
- In older versions, set **`sched.cpu.latencySensitivity`** to High.

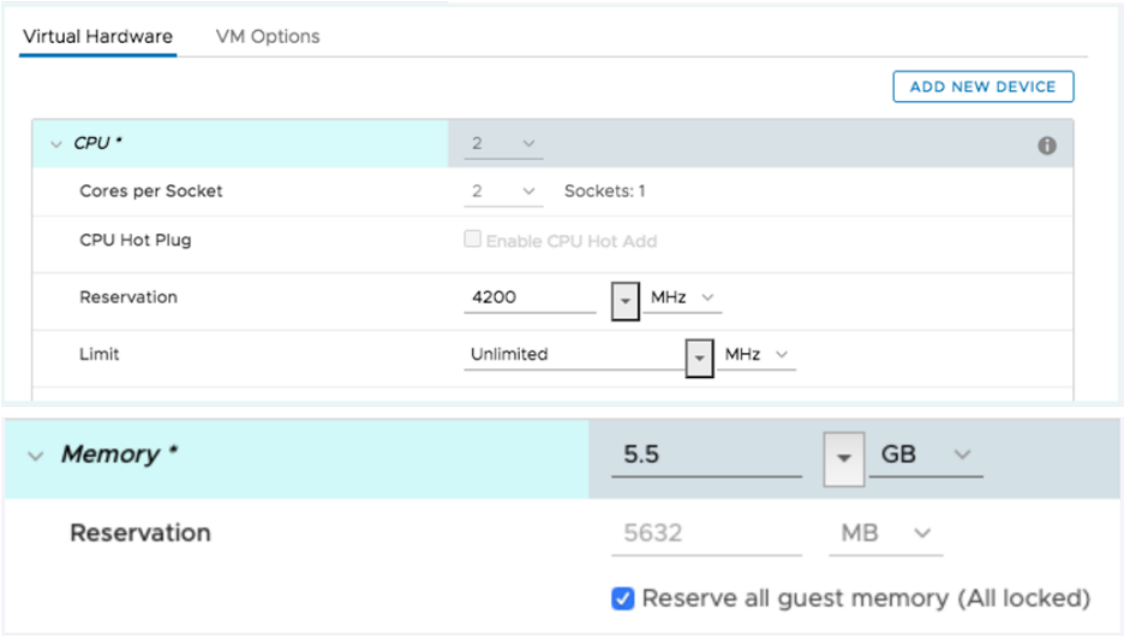

 ©2024 Palo Alto Networks, Inc.

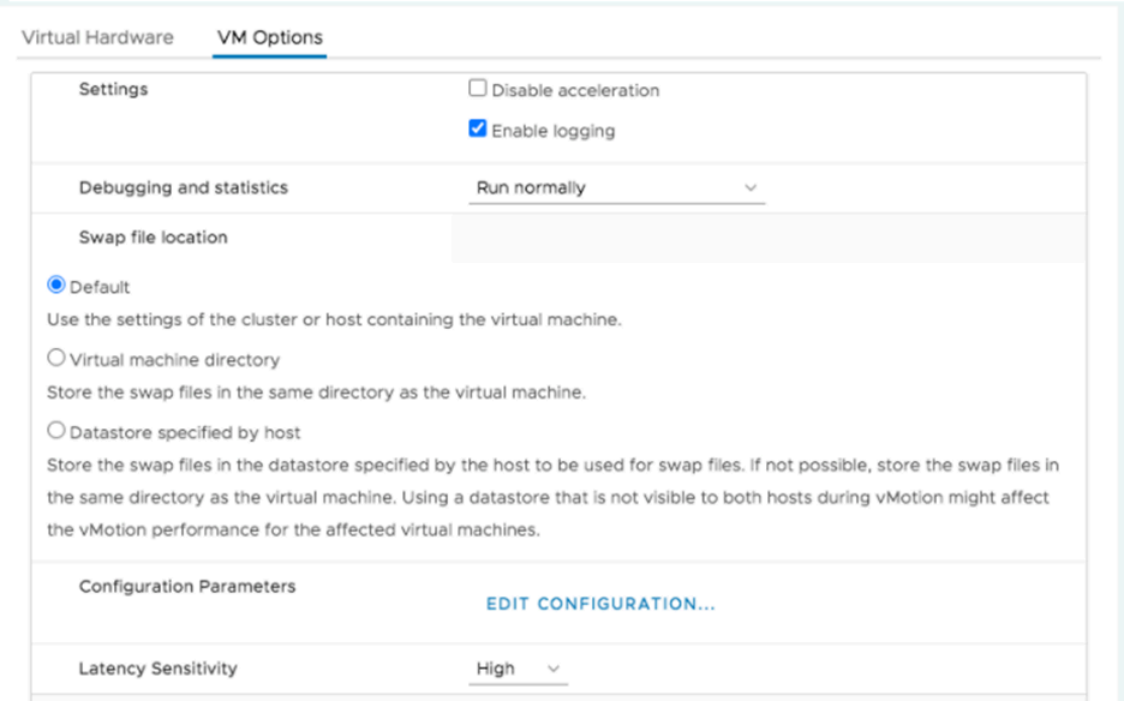

Additionally, VM's vCPUs can be pinned to host CPU cores using the VM setting **Host Affinity** so that it is never scheduled to different cores. Keep NUMA and hyperthreading in mind when you use Host Affinity. Avoid setting Host Affinity if the system is over committed. For more detail see Potential Issues with CPU Affinity.

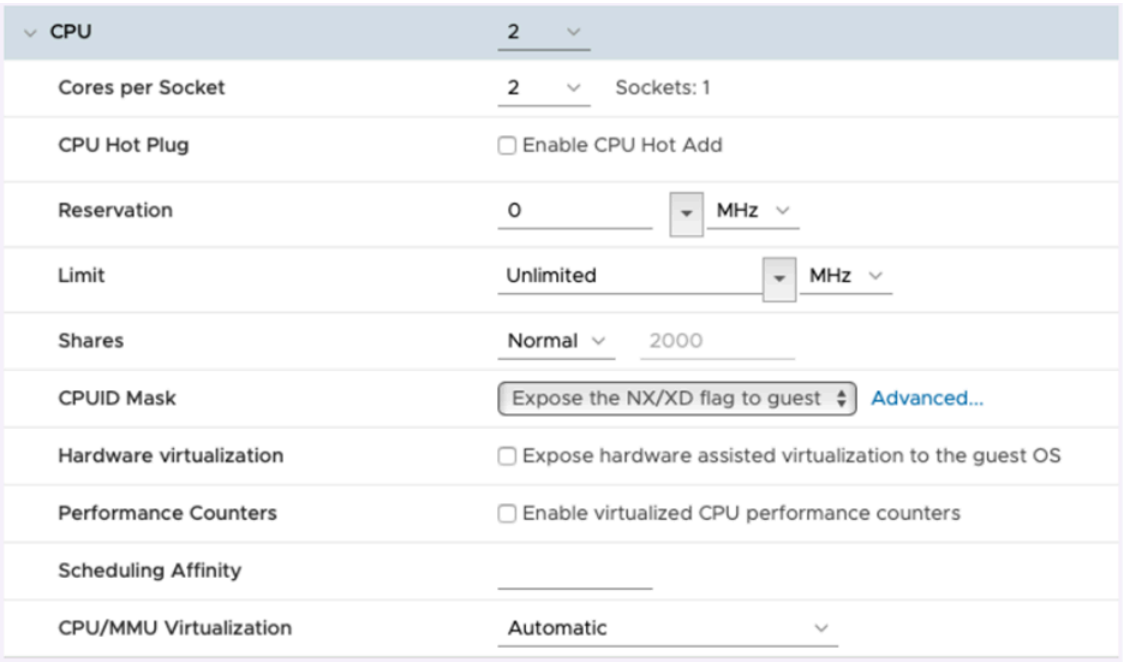

After you implement the tuning parameters, use esxtop or CPU charts to check CPU Ready (%RDY) and Co Stop (%CSTP) for the VM. Both values should be close to 0% to ensure exclusive access to CPU resources. You can also use esxtop to check for NUMA usage and ensure memory resources for the VM are not spread across NUMA nodes. For more detail, see Interpreting esxtop Statistics.

 ©2024 Palo Alto Networks, Inc.

Use Cases

Use Case 1: vSwitch Deployment

The figure below shows a deployment of a PA-VM on an ESXi host where the data ports “Port 1” and “Port 2” are linked to eth1 and eth2 of the PA-VM. Each port hosts two queue pairs (for example, Tx0/Rx0, and Tx1/Rx1) or has multiqueue enabled.

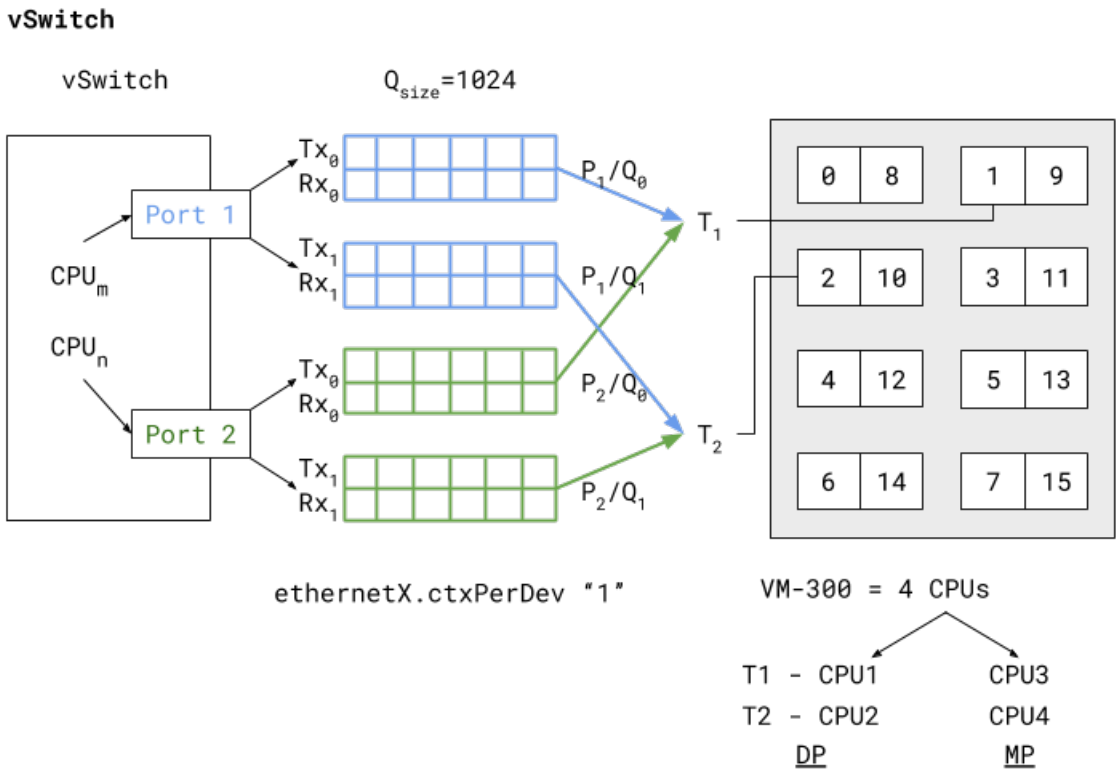

Enabling multiqueue and RSS for load balancing packets sent/received to/from multiple queues enhances processing performance. Based on an internal logic of vCPU to port/queue mapping (in this case) packets arriving and being sent out from P1/Q0 and P2/Q0 are processed by dataplane task T1 running on (i.e., pinned to) vCPU1. The data plane task T2 follows a similar association, as shown in the vSwitch deployment diagram above.

The two data plane tasks are running on vCPU1 and vCPU2 and these are non-sibling CPUs (means that they do not share the same core in case of hyperthreading). This means that even with hyperthreading enabled the task assignment can be pinned to different cores for high performance. Also these dataplane task vCPUs all belong to the same NUMA node (or socket) to avoid NUMA-related performance issues.

Two other performance bottlenecks can be addressed with increasing the queue sizes and dedicating a vCPU or thread to the ports that schedule traffic to and from these ports. Increasing the queue sizes (Qsize) will accommodate large sudden bursts of traffic and prevent packet drops under bursty traffic. Adding a dedicated CPU thread (**`ethernetX.ctxPerDev = 1`**) to port level packet processing will allow traffic to be processed at a higher rate, thereby increasing the traffic throughput to reach line rate.

The PA-VM packet processing technique also determines performance. This can be set to either DPDK or PacketMMAP. DPDK uses a poll mode driver (depends on the driver type) to constantly poll for packets received in the queues. This leads to higher throughput performance. Depending

 ©2024 Palo Alto Networks, Inc.

on the poll period is latency observed by the packets. If the polling is continuous (i.e., busy-poll a setting from the PANOS cli) then the vCPU utilization for the data plane tasks will be a 100% but will yield the best performance. Internally the software uses a millisecond-level polling time to prevent unnecessary utilization of CPU resources.

PacketMMAP, on the other hand, has a lower performance than DPDK but it works with any network level drivers. For DPDK the vSwitch driver must have support for DPDK. PacketMMAP works with interrupts that are raised when a packet is received by the port and placed in the receive queue. This means that for every packet, or group of packets, interrupts are raised and packets are drained off the receive queue for processing. This results in lower latency in packet processing, but reduced throughput, because interrupts must be processed every time, causing higher CPU overhead. In general PacketMMAP will have lower packet processing latency than DPDK (without busy poll modification).

Use Case 2: SR-IOV Deployment

The SR-IOV diagram below shows a PAVM deployment similar to the vSwitch use case, but in SR-IOV mode.

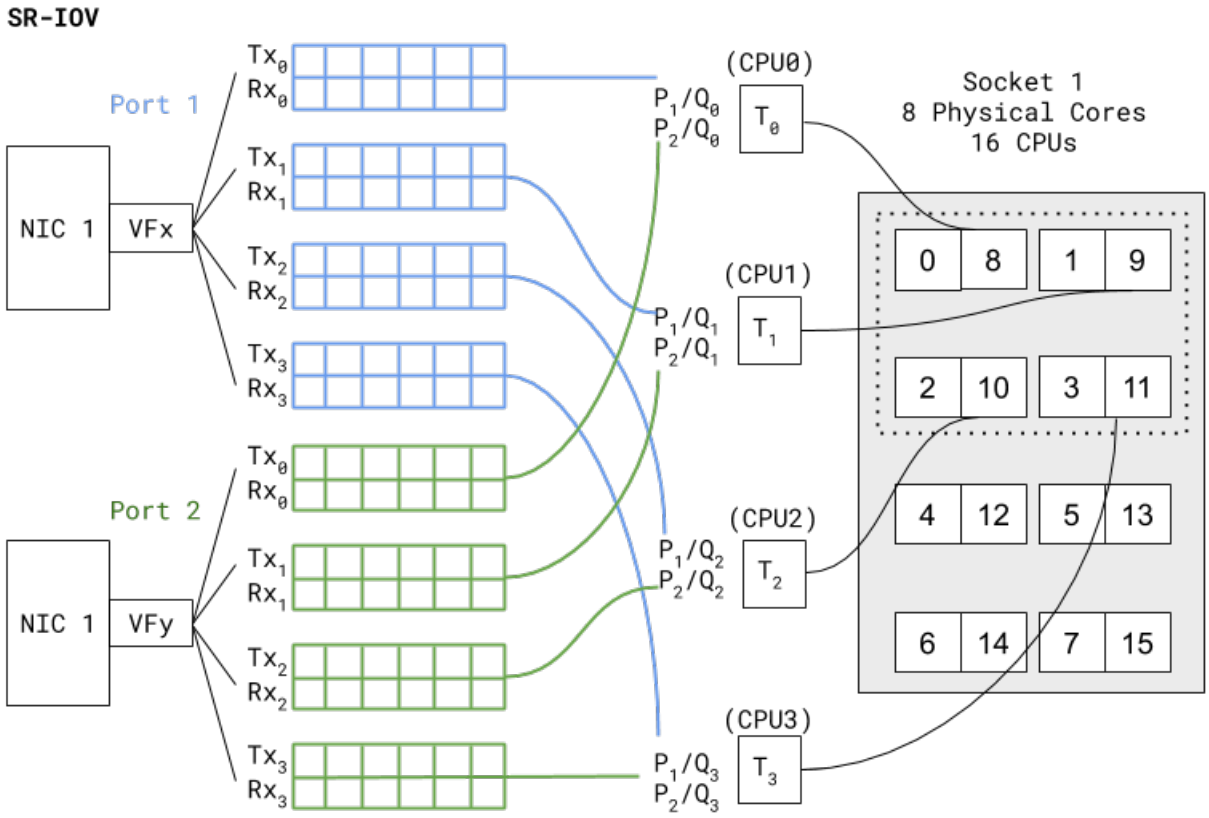

In SR-IOV the compatible physical NIC port (manifests as a Physical Function) is essentially carved out into multiple interfaces (manifests as Virtual Functions). The figure above shows that NIC1 Port1 has a VF named VFX that is associated as one of the PAVM dataplane interfaces — eth1, for example. A similar association is created for Port2 VF to PAVM eth2.The chain of packet processing is similar to that of the deployment in the vSwitch environment. The only difference is that the SR-IOV VF drivers should be compatible with those used in PAN-OS. Also, since there is no internal vSwitch (in the host) switching traffic, there is no need to set a dedicated thread for traffic scheduling from a port (that is, **`ethernetX.ctxPerDev = 1`** is not required in this

 ©2024 Palo Alto Networks, Inc.

setting). Interfaces with SR-IOV and DPDK will yield even higher packet processing performance than the vSwitch use case.

References

- Tuning VMware vCloud NFV for Data-Intensive Workloads
- Best Practices for Performance Tuning of Telco and NFV Workloads in vSphere
- Potential Issues with CPU Affinity
- Interpreting esxtop Statistics

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on vCloud Air

The VM-Series firewall can be deployed in a virtual data center (vDC) on vCloud Air using the vCloud Air portal, from the vCloud Director portal or using the vCloud Air API.

- About the VM-Series Firewall on vCloud Air
- Deployments Supported on vCloud Air
- Deploy the VM-Series Firewall on vCloud Air

About the VM-Series Firewall on vCloud Air

You can deploy the VM-Series firewall in a virtual data center (vDC) on VMware vCloud Air using the vCloud Air portal or from the vCloud Director portal. And to centrally manage all your physical and VM-Series firewalls, you can use an existing Panorama or deploy a new Panorama on premise or on vCloud Air.

The VM-Series firewall on vCloud Air requires the following:

- ESXi version of the software image, an Open Virtualization Alliance (OVA) file, from the Palo Alto Networks Customer Support web site. Currently, the vCloud Air Marketplace does not host the software image.

  In order to efficiently deploy the VM-Series firewall, include the firewall software image in a vApp. A vApp is a container for preconfigured virtual appliances (virtual machines and operating system images) that is managed as a single object. For example, if your vApp includes a set of multi-tiered applications and the VM-Series firewall, each time you deploy the vApp, the VM-Series firewall automatically secures the web server and database server that get deployed with the vApp.
- License and subscriptions purchased from a partner, reseller, or directly from Palo Alto Networks, in the Bring Your Own License (BYOL) model; the usage-based licensing for the VM-Series on vCloud Air is not available.
- Due to the security restrictions imposed on vCloud Air, the VM-Series firewall on vCloud Air is best deployed with Layer 3 interfaces and the interfaces must be enabled to use the hypervisor assigned MAC address. If you do not enable hypervisor assigned MAC address, the VMware vSwitch cannot forward traffic to the dataplane interfaces on the VM-Series firewall because the vSwitch on vCloud Air does not support promiscuous mode or MAC forged transmits. The VM-Series firewall cannot be deployed with tap interfaces, Layer 2 interfaces, or virtual wire interfaces.

The VM-Series firewall on vCloud Air can be deployed in an active/passive high availability configuration. However, the VM-Series firewall on vCloud Air does not support VM Monitoring capabilities for virtual machines that are hosted on vCloud Air.

To learn all about vCloud Air, refer to the VMware vCloud Air documentation.

 ©2024 Palo Alto Networks, Inc.

Deployments Supported on vCloud Air

To enable applications safely, block known and unknown threats, and to keep pace with changes in your environment, you can deploy the VM-Series firewall on vCloud Air with Layer 3 interfaces in the following ways:

- **Secure the virtual data center perimeter**—Deploy the VM-Series firewall as a virtual machine that connects isolated and routed networks on vCloud Air. In this deployment the firewall secures all north-south traffic traversing the infrastructure on vCloud Air.
- **Set up a hybrid cloud**—Extend your data center and private cloud into vCloud Air and use a VPN connection to enable communication between the corporate network and the data center. In this deployment, the VM-Series firewall uses IPSec to encrypt traffic and secure users accessing the cloud.
- **Secure traffic between application subnets in the vDC**—To improve security, segment your network and isolate traffic by creating application tiers, and then deploy the VM-Series firewall to protect against lateral threats between subnets and application tiers.

The following illustration combines all three deployments scenarios and includes Panorama. Panorama streamlines policy updates, centralizes policy management, and provides centralized logging and reporting.

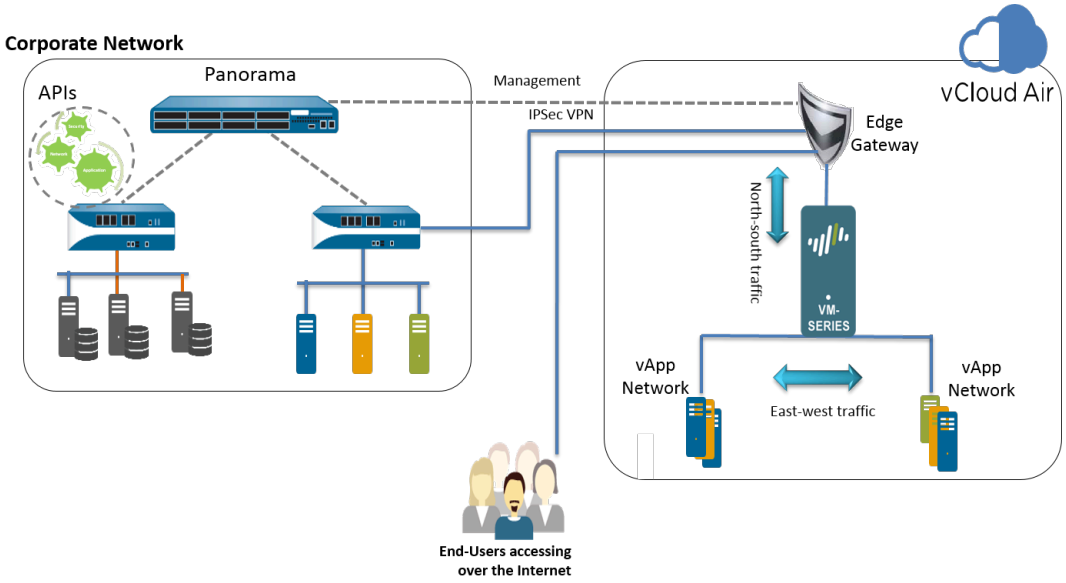

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on vCloud Air

Use the instructions in this section to deploy your VM-Series firewall in an on-demand or dedicated vDC on vCloud Air. This procedure assumes that you have set up your vDC, including the gateways required to allow traffic in and out of the vDC, and the networks required for routing management traffic and data traffic through the vDC.

STEP 1 | Obtain the VM-Series OVA image from the Palo Alto Networks Customer Support web site; the vCloud Air Marketplace does not host the software image currently.

1. Go to: www.paloaltonetworks.com/services/support.html.
2. Filter by **PAN-OS for VM-Series Base Images** and download the OVA image. For example, PA-VM-ESX-9.1.0.ova.

STEP 2 | Extract the Open Virtualization Format (OVF) file from the OVA image and import the OVF file in to your vCloud Air catalog.

When extracting files from the OVA image, make sure to place all the files—.mf, .ovf, and .vmdk—within the same directory.

For instructions to extract the OVF file from the OVA image, refer to the VMware documentation: https://www.vmware.com/support/developer/ovf/#sthash.WUp55ZyE.dpuf

When you import the OVF file, the software image for the VM-Series firewall is listed in **My Organization's Catalogs**.

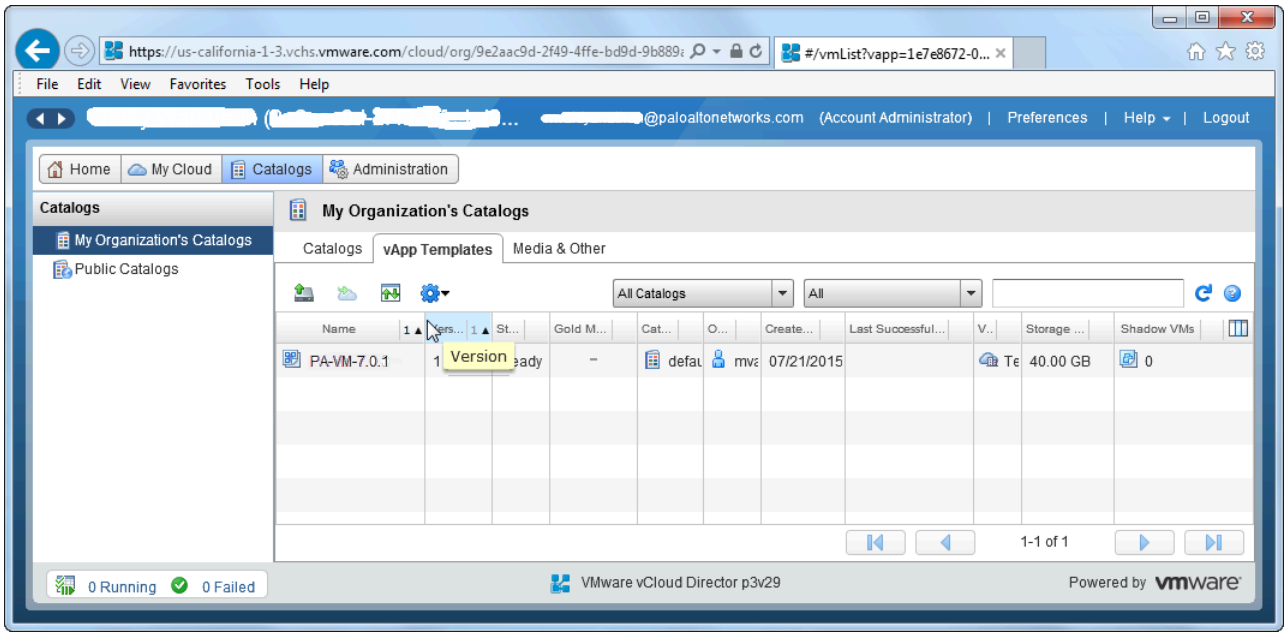

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Choose your workflow.

A vApp is a collection of templates for preconfigured virtual appliances that contain virtual machines, and operating system images.

- If you want to create a new vDC and a new vApp that includes the VM-Series firewall, go to step 4
- If you have already deployed a vDC and have a vApp and now want to add the VM-Series firewall to the vApp to secure traffic, go to step 5

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Create a vDC and a vApp that includes the VM-Series firewall.

1. Log in to vCloud Air.
2. Select **VPC OnDemand** and select the location in which you want to deploy the VM-Series firewall.

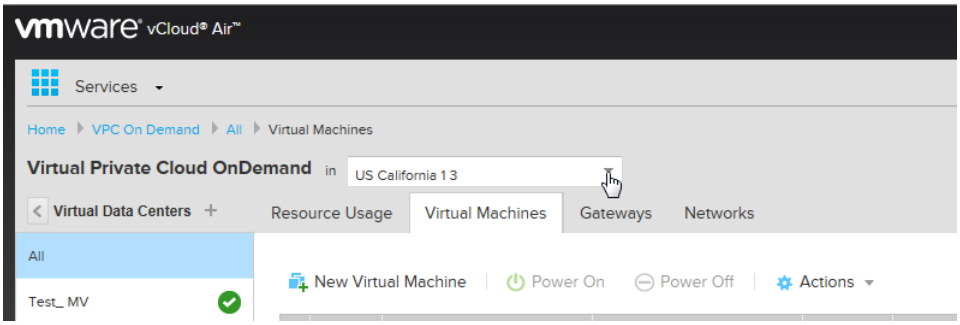

3. Select **Virtual Data Centers** and click **+** to add a new Virtual Data Center.
4. Select the vDC, right click and select **Manage Catalogs in vCloud Director**. You will be redirected to the vCloud Director web interface.
5. Create a new vApp that contains one or more virtual machines including the VM-Series firewall:
   1. Select **My Cloud** > **vApps**, and click **Build New vApp**.

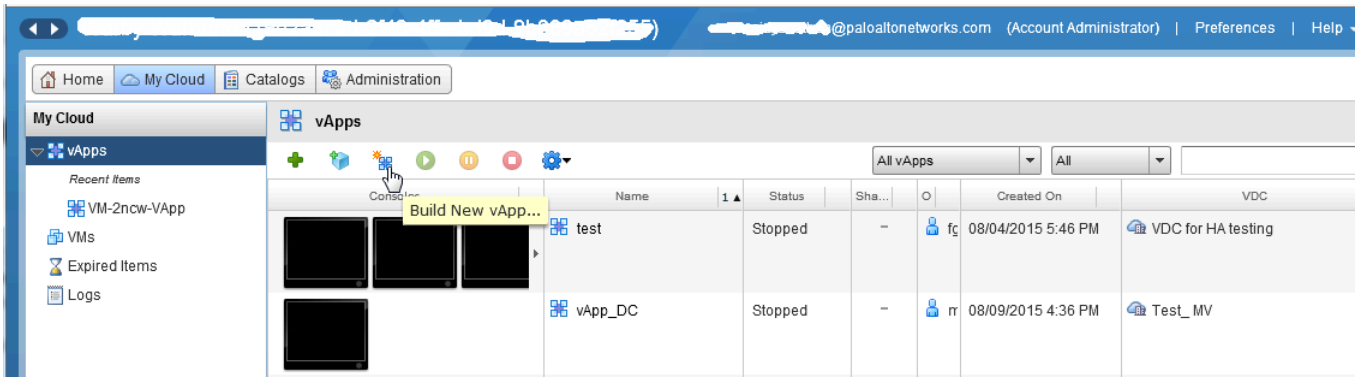

   2. Select **Name and Location**, and the **Virtual Datacenter** in which this vApp will run. By default, **Leases** for runtime and storage never expire and the vApp is not automatically stopped.
   3. **Add Virtual Machines**. To add the VM-Series firewall image from the **Look in:** drop-down, select **My Organization's Catalog**, select the image and click **Add**. Click **Next**
   4. Configure **Resources** to specify the Storage Policies for the virtual machines when deployed. The VM-Series firewall uses the **Standard** option.
   5. Configure the **Virtual Machines**. Name each virtual machine and select the network to which you want it to connect. You must connect NIC 0 (for management access) to the default routed network; NIC 1 is used for data traffic. You can add additional NICs later.
   6. Verify the settings and click **Finish.**
   7. Continue to step 6.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Add the VM-Series Firewall into a vApp.

1. Log in to vCloud Air.
2. Select your existing **Virtual Data Center** from the left pane, right click and select **Manage Catalogs in vCloud Director**. You will be redirected to the vCloud Director web interface.
3. Select **My Cloud** > **vApps** and click the **Name** of the vApp in which to include the VM-Series firewall.
4. Open the vApp (double-click on the name), select **Virtual Machines** and click ✚ to add a virtual machine.
    1. In the **Look in:** drop-down, choose **My Organization's Catalog**, select the VM-Series firewall image and click **Add**. Click **Next**.
    2. Click **Next to** skip **Configure Resources**. The VM-Series firewall uses the **Standard** option and you do not to modify the Storage Policy.
    3. Enter a **Name** for the firewall and for management access **(NIC 0)**, select the default routed network and the **IP Mode**— Static or DHCP. You can configure NIC 1 and add additional NICs in step 6. Click **Next**.
    4. Verify how this vApp connects to the vDC— Gateway Address and Network Mask for the virtual machines in this vApp.
    5. Verify that you have added the VM-Series firewall and click **Finish.**
    6. Continue to step 6.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Connect the data interface(s) of the VM-Series firewall to an isolated or a routed network, as required for your deployment.

1. In vCloud Director, select **My Cloud** > **vApps** and select the vApp you just created or edited.
2. Select **Virtual Machines** and select the VM-Series firewall. Then, right-click and select **Properties**.
3. Select **Hardware**, scroll to the NICs section and select **NIC 1**.
4. Attach the dataplane network interface to a vApp network or an organizational VDC network based on your connectivity needs for data traffic to the VM-Series firewall. To create a new network:
   1. In the Network drop-down, click **Add Network**.
   2. Select the **Network Type** and give it a name and click **OK**.
   3. Verify that the new network is attached to the interface.
5. To add additional NICs to the firewall, click **Add** and repeat step 4 above. You can attach a maximum of seven dataplane interfaces to the VM-Series firewall.
6. Verify that the management interface of the VM-Series firewall is attached to the default routed subnet on the vDC and at least one dataplane interface is connected to a routed or isolated network.
   1. Select **My Cloud** > **vApps** and double-click the **Name** of the vApp you just edited.
   2. Verify network connectivity in the **vApp Diagram**.

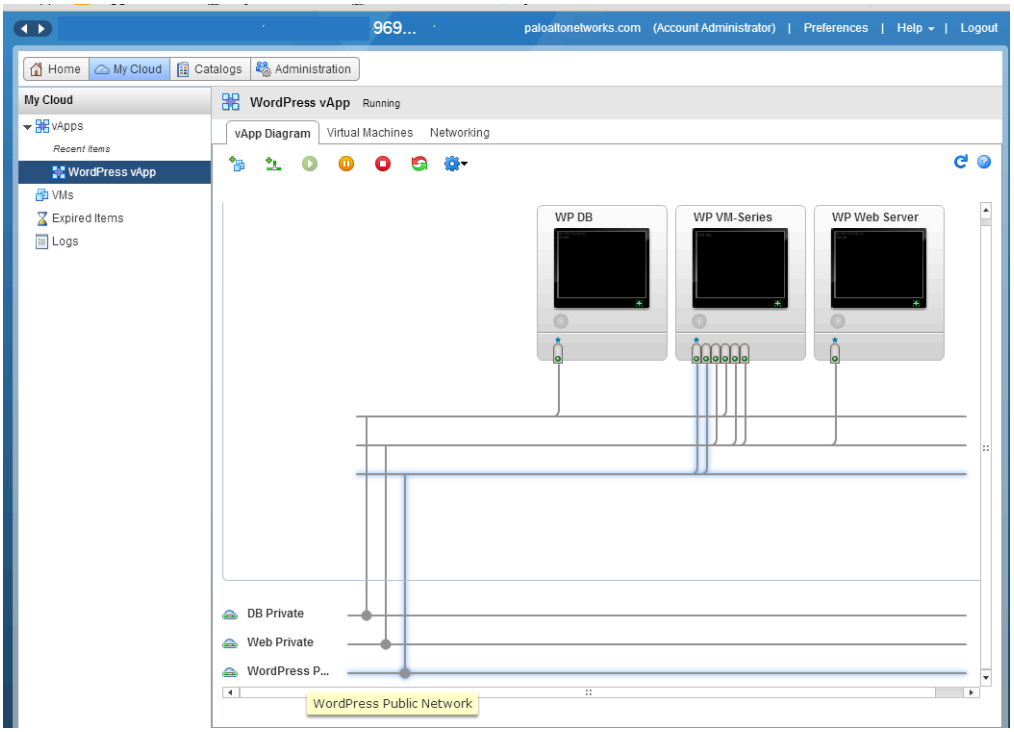

 ©2024 Palo Alto Networks, Inc.

STEP 7 | (Optional) Edit the hardware resources allocated for the VM-Series firewall.

Required only if you need to allot additional CPU, memory, or hard disk to the firewall.

1. Select **My Cloud** > **vApps** and double-click the **Name** of the vApp you just deployed.

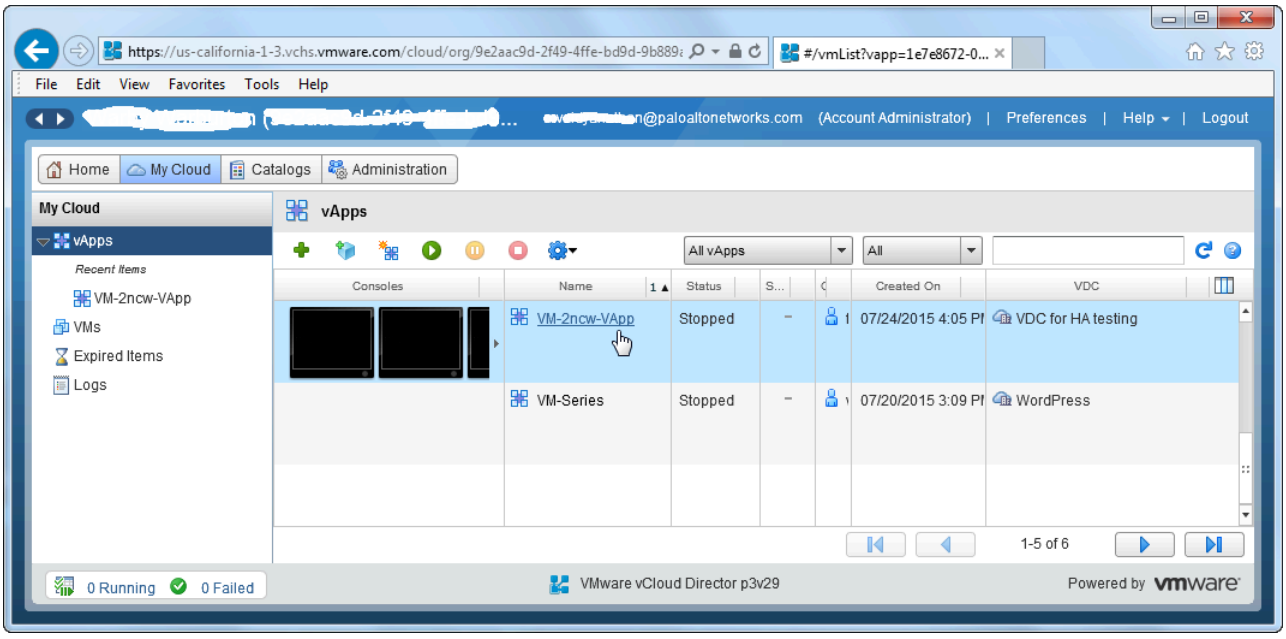

2. Select **Virtual Machine** and click on the **Name** of the VM-Series firewall to access the Virtual Machine Properties.

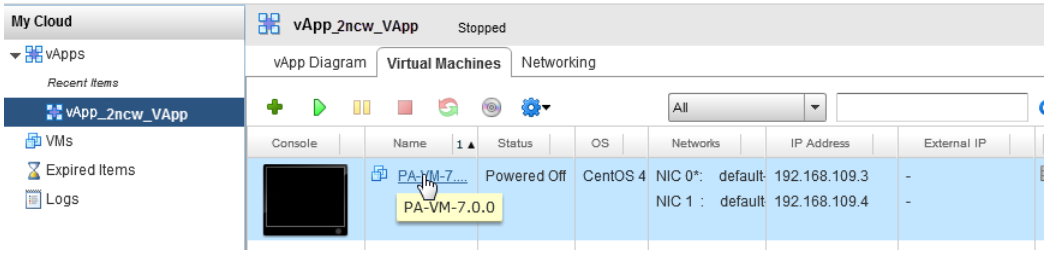

3. Add additional **Hardware** resources for the VM-Series firewall:
   - See VM-Series System Requirements for the minimum vCPU, memory, and disk requirements for your VM-Series model.
   - NICs: One management and up to seven dataplane interfaces.

STEP 8 | Power on the VM-Series firewall.

STEP 9 | Configure an IP address for the VM-Series firewall management interface.

Perform Initial Configuration on the VM-Series on ESXi.

The VM-Series firewall on vCloud Air supports VMware Tools, and you can Use VMware Tools on the VM-Series Firewall on ESXi and vCloud Air to view the management IP address of the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Define NAT rules on the vCloud Air Edge Gateway to enable Internet access for the VM-Series firewall.

1. Select **Virtual Data Centers** > **Gateways**, select the gateway and double-click to add **NAT Rules**.
2. Create two DNAT rules. One for allowing SSH access and one for HTTPS access to the management port's IP address on the VM-Series firewall.
3. Create a SNAT rule for translating the internal source IP address for all traffic initiated from the management port on the VM-Series firewall to an external IP address.

To send and receive traffic from the dataplane interfaces on the firewall, you must create additional DNAT and SNAT rules on the vCloud Air Edge Gateway.

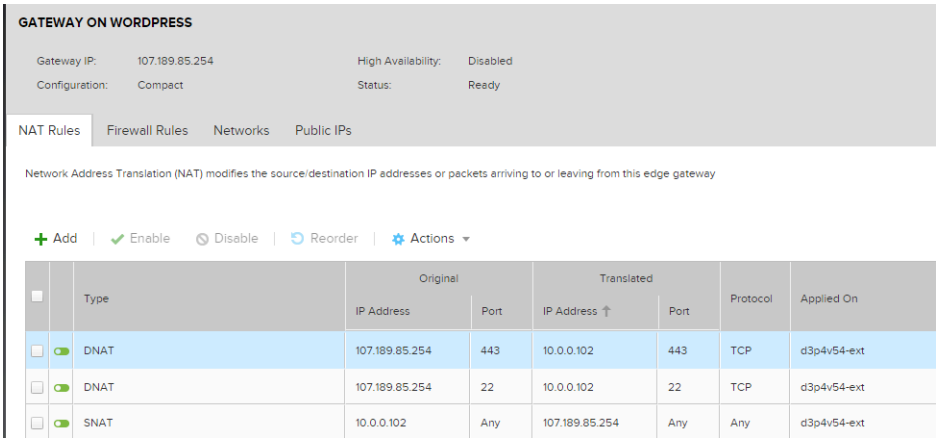

STEP 11 | Log in to the web interface of the firewall.

In this example, the URL for the web interface is https://107.189.85.254

The NAT rule on the Edge Gateway translates the external IP address and port 107.189.85.254:443 to the private IP address and port 10.0.0.102:443.

STEP 12 | Add the auth code(s) to activate the licenses on the firewall.

Activate the License.

STEP 13 | Configure the VM-Series firewall to use the hypervisor assigned MAC address.

Hypervisor Assigned MAC Addresses

 ©2024 Palo Alto Networks, Inc.

STEP 14 | Configure the dataplane interfaces as Layer 3 interfaces.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type: Layer3**
   - Select the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. Define a new zone, for example untrust, and then click **OK**.
   - Select **IPv4**, assign a static IP address.
   - On **Advanced** > **Other Info**, expand the **Management Profile** drop-down, and select **New Management Profile**.
   - Enter a **Name** for the profile, such as allow_ping, and select Ping from the Permitted Services list, then click **OK**.
   - To save the interface configuration, click **OK**.
3. Repeat the process for each additional interface.
4. Click **Commit** to save the changes.

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on VMware NSX

The VM-Series firewall can be deployed in both versions of VMware's network virtualization solution—NSX-V and NSX-T.

- Set Up the VM-Series Firewall on VMware NSX-V
- Set Up the VM-Series Firewall on VMware NSX-T (North-South)
- Set Up the VM-Series Firewall on NSX-T (East-West)

Set Up the VM-Series Firewall on VMware NSX-V

The VM-Series firewall for VMware NSX-V is jointly developed by Palo Alto Networks and VMware. This solution uses the NetX API to integrate the Palo Alto Networks next-generation firewalls and Panorama with VMware ESXi servers to provide comprehensive visibility and safe application enablement of all data center traffic including intra-host virtual machine communications.

The following topics provide information about the VM-Series for NSX-V:

- VM-Series for Firewall NSX-V Overview
- VM-Series Firewall for NSX-V Deployment Checklist
- Install the VMware NSX Plugin
- Register the VM-Series Firewall as a Service on the NSX-V Manager
- Deploy the VM-Series Firewall
- Create Security Groups and Steering Rules
- Apply Security Policies to the VM-Series Firewall
- Steer Traffic from Guests that are not Running VMware Tools
- What is Multi-NSX Manager Support on the VM-Series for NSX-V?
- Add a New Host to Your NSX-V Deployment
- Dynamically Quarantine Infected Guests
- Migrate Operations-Centric Configuration to Security-Centric Configuration
- Use Case: Shared Compute Infrastructure and Shared Security Policies
- Use Case: Shared Security Policies on Dedicated Compute Infrastructure
- Dynamic Address Groups—Information Relay from NSX-V Manager to Panorama

VM-Series for Firewall NSX-V Overview

NSX-V, VMware's Networking and Security platform designed for the software-defined data center (SDDC), offers the ability to deploy the Palo Alto Networks firewall as a service on a cluster of ESXi servers. The term *SDDC* is a VMware term that refers to a data center where infrastructure—compute resources, network and storage—is virtualized using VMware NSX-V.

To keep pace with the changes in the agile SDDC, the VM-Series firewall for NSX-V simplifies the process of deploying a Palo Alto Networks next-generation firewall and continually enforcing security and compliance for the east-west traffic in the SDDC. For details on the VM-Series for NSX-V, see the following topics:

- What are the Components of the VM-Series for NSX-V Solution?
- How Do the Components in the VM-Series Firewall for NSX-V Solution Work Together?
- What are the Benefits of the NSX-V VM-Series firewall for NSX-V Solution?
- What is Multi-Tenant Support on the VM-Series Firewall for NSX-V?

 ©2024 Palo Alto Networks, Inc.

What are the Components of the VM-Series for NSX-V Solution?

The following tables show the components of this joint Palo Alto Networks and VMware solution. The following topics describe each component in more detail:

- vCenter Server
- NSX-V Manager
- Panorama
- VM-Series Firewall for NSX-V
- Ports/Protocols used Network Communication

VMware Components	
vCenter Server	The vCenter server is the centralized management tool for the vSphere suite.
NSX-V Manager	VMware's Networking and Security platform must be installed and registered with the vCenter server. The NSX-V Manager is required to deploy the VM-Series firewall on the ESXi hosts within a ESXi cluster.
ESXi Server	ESXi is a hypervisor that enables compute virtualization.

See the Palo Alto Networks Compatibility Matrix for supported software versions.

Palo Alto Networks Components	
PAN-OS	The VM-Series base image (PA-VM-NSX-9.1.zip) is used for deploying the VM-Series firewall for NSX-V with PAN-OS 9.1. The minimum system requirement for deploying the VM-Series firewall for NSX-V on the ESXi server depends on your VM-Series model. See VM-Series System Requirements for the minimum hardware requirements for your VM-Series model.
Panorama Panorama must be running the same release version or later version that the firewalls that it will manage.	Panorama is the centralized management tool for the Palo Alto Networks next-generation firewalls. In this solution, Panorama works with the NSX-V Manager to deploy, license, and centrally administer—configuration and policies—on the VM-Series firewall for NSX-V. Panorama must be able to connect to the NSX-V Manager, the vCenter server, the VM-Series firewalls and the Palo Alto Networks update server. The resources required by Panorama depend on the mode Panorama will run in: Panorama or Management Only. Panorama installations running in Legacy mode prior to upgrade remain in Legacy mode after upgrading to 10.0.

 ©2024 Palo Alto Networks, Inc.

Palo Alto Networks Components	
	See the Panorama Administrator's Guide for information about deploying your Panorama appliance.
VM-Series Firewall for NSX-V	The VM-100, VM-200, VM-300, VM-500, and VM-1000-HV, support NSX-V.

vCenter Server

The vCenter server is required to manage the NSX-V Manager and the ESXi hosts in your data center. This joint solution requires that the ESXi hosts be organized into one or more clusters on the vCenter server and must be connected to a distributed virtual switch.

For information on clusters, distributed virtual switch, DRS, and the vCenter server, refer to your VMware documentation: http://www.vmware.com/support/vcenter-server.html

NSX-V Manager

NSX-V is VMware's network virtualization platform that is completely integrated with vSphere. The NSX-V Firewall and the Service Composer are key features of the NSX-V Manager. The NSX-V firewall is a logical firewall that allows you to attach network and security services to the virtual machines, and the Service Composer allows you to group virtual machines and create policy to redirect traffic to the VM-Series firewall (called the Palo Alto Networks NGFW service on the NSX-V Manager).

Panorama

Panorama is used to register the VM-Series firewall for NSX-V as the *Palo Alto Networks NGFW* service on the NSX-V Manager. Registering the Palo Alto Networks NGFW service on the NSX-V Manager allows the NSX-V Manager to deploy the VM-Series firewall for NSX-V on each ESXi host in the ESXi cluster.

Panorama serves as the central point of administration for the VM-Series firewalls running on NSX-V. When a new VM-Series firewall is deployed in NSX-V, it communicates with Panorama to obtain the license and receives its configuration/policies from Panorama. All configuration elements, policies, and dynamic address groups on the VM-Series firewalls can be centrally managed on Panorama using Device Groups and Template Stacks. The REST-based XML API integration in this solution, enables Panorama to synchronize with the NSX-V Manager and the VM-Series firewalls to allow the use of dynamic address groups and share context between the virtualized environment and security enforcement. For more information, see Policy Enforcement using Dynamic Address Groups.

VM-Series Firewall for NSX-V

The VM-Series firewall for NSX-V is the VM-Series firewall that is deployed on the ESXi hypervisor. The integration with the NetX API makes it possible to automate the process of installing the VM-Series firewall directly on the ESXi hypervisor, and allows the hypervisor to forward traffic to the VM-Series firewall without using the vSwitch configuration; it therefore, requires no change to the virtual network topology.

The VM-Series firewall for NSX-V only supports virtual wire interfaces. On this firewall, ethernet 1/1 and ethernet 1/2 are bound together through a virtual wire and use the NetX dataplane API to communicate with the hypervisor. Layer 2 or Layer 3 interfaces are neither

 ©2024 Palo Alto Networks, Inc.

required nor supported on the VM-Series firewall for NSX-V, and therefore no switching or routing actions can be performed by the firewall. For enabling traffic separation in a multi-tenancy environment, you can create additional zones that internally map to a pair of virtual wire subinterfaces on the parent virtual wire interfaces, ethernet 1/1 and ethernet 1/2.

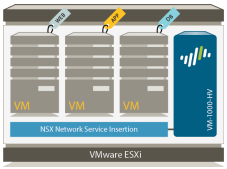

Ports/Protocols used Network Communication

In order to enable the network communication required to deploy the VM-Series firewall for NSX-V, you must allow the use of the following protocols/ports and applications.

- **Panorama**—To obtain software updates and dynamic updates, Panorama uses SSL to access updates.paloaltonetworks.com on TCP/443; this URL leverages the CDN infrastructure. If you need a single IP address, use staticupdates.paloaltonetworks.com. The App-ID for updates is paloalto-updates.

  The NSX-V Manager and Panorama use SSL to communicate on TCP/443.
- **VM-Series Firewall for NSX-V**—If you plan to use WildFire, the VM-Series firewalls must be able to access wildfire.paloaltonetworks.com on port 443. This is an SSL connection and the App-ID is paloalto-wildfire-cloud.

  The management interface on the VM-Series firewall uses SSL to communicate with Panorama over TCP/3978.
- **vCenter Server** The vCenter Server must be able to reach the deployment web server that is hosting the VM-Series OVA. The port is TCP/80 by default or App-ID web-browsing.

How Do the Components in the VM-Series Firewall for NSX-V Solution Work Together?

To meet the security challenges in the software-defined data center, the NSX-V Manager, ESXi servers and Panorama work harmoniously to automate the deployment of the VM-Series firewall.

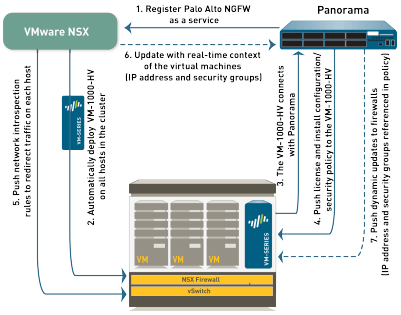

1. **Register the Palo Alto Networks NGFW service**—The first step is to register the Palo Alto Networks NGFW as a service on the NSX-V Manager. The registration process uses the NetX management plane API to enable bi-directional communication between Panorama and the NSX-V Manager. Panorama is configured with the IP address and access credentials to initiate a connection and register the Palo Alto Networks NGFW service on the NSX-V Manager. The service definition includes the URL for accessing the VM-Series base image that is required to deploy the VM-Series firewall for NSX-V, the authorization code for retrieving the license and

 ©2024 Palo Alto Networks, Inc.

the device group and template stacks to which the VM-Series firewalls will belong. The NSX-V manager uses this management plane connection to share updates on the changes in the virtual environment with Panorama.

2. **Deploy the VM-Series automatically from NSX-V**—The NSX-V Manager collects the VM-Series base image from the URL specified during registration and installs an instance of the VM-Series firewall on each ESXi host in the ESXi cluster. From a static management IP pool or a DHCP service (that you define on the NSX-V Manager), a management IP address is assigned to the VM-Series firewall and the Panorama IP address is provided to the firewall. When the firewall boots up, the NetX dataplane integration API connects the VM-Series firewall to the hypervisor so that it can receive traffic from the vSwitch.

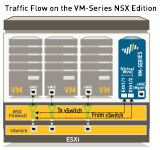

3. **Establish communication between the VM-Series firewall and Panorama**—The VM-Series firewall then initiates a connection to Panorama to obtain its license. Panorama retrieves the license from the update server and pushes it to the firewall. The VM-Series firewall receives the license and reboots with a valid serial number.

   If your Panorama is offline, which means that it does not have direct Internet access to retrieve the licenses and push them to the firewalls, you must manually license each firewall. If your VM-Series firewall does not have internet access, you must add the serial number of the firewall to Panorama so that it is registered as a managed device, so that you can push the appropriate template stacks and device group settings from Panorama.

4. **Install configuration/policy from Panorama to the VM-Series firewall**—The VM-Series firewall reconnects with Panorama and provides its serial number. Panorama now adds the firewall to the device group and template stack that was defined in the service definition and pushes the configuration and policy rules to the firewall. The VM-Series firewall is now available as a security virtual machine that can be further configured to safely enable applications on the network.
5. **Push traffic redirection rules to NSX-V Manager**—Create security groups and define network introspection rules that specify the guests from which traffic will be steered to the VM-Series firewall. See Integrated Policy Rules for details.

   To ensure that traffic from the guests is steered to the VM-Series firewall, you must have VMware Tools installed on each guest. If VMware Tools is not installed, the NSX-V Manager does not know the IP address of the guest and therefore, the traffic cannot be steered to the VM-Series firewall. For more information, see Steer Traffic from Guests that are not Running VMware Tools. This is not required if you are running NSX-V Manager 6.2.4 or later.

6. **Receive real-time updates from NSX-V Manager**—The NSX-V Manager sends real-time updates on the changes in the virtual environment to Panorama. These updates include information on the security groups and IP addresses of guests that are part of the security group from which traffic is redirected to the VM-Series firewall. See Integrated Policy Rules for details.

 ©2024 Palo Alto Networks, Inc.

7. **Use dynamic address groups in policy and push dynamic updates from Panorama to the VM-Series firewalls**—On Panorama, use the real-time updates on security groups to create dynamic address groups, bind them to security policies and then push these policies to the VM-Series firewalls. Every VM-Series firewall in the device group will have the same set of policies and is now completely marshaled to secure the SDDC. See Policy Enforcement using Dynamic Address Groups for details.

Integrated Policy Rules

Panorama serves as the single point of configuration that provides the NSX-V Manager with the contextual information required to redirect traffic from the guest virtual machines to the VM-Series firewall. The traffic steering rules are defined on Panorama and pushed to NSX-V Manager; these determine what traffic from which guests in the cluster are steered to the Palo Alto Networks NGFW service. Security enforcement rules are also defined on Panorama and pushed to the VM-Series firewalls for the traffic that is steered to the Palo Alto Networks NGFW service.

- **Steering Rules**—The rules for directing traffic from the guests on each ESXi host are defined on Panorama and applied by NSX-V Manager as partner security services rules.

  For traffic that needs to be inspected and secured by the VM-Series firewall, the steering rules created on Panorama allow you to redirect the traffic to the Palo Alto Networks NGFW service. This traffic is then steered to the VM-Series firewall and is first processed by the VM-Series firewall before it goes to the virtual switch.

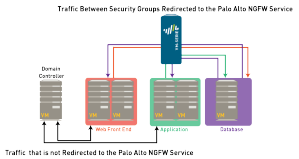

  Traffic that does not need to be inspected by the VM-Series firewall, for example network data backup or traffic to an internal domain controller, does not need to be redirected to the VM-Series firewall and can be sent to the virtual switch for onward processing.

- **Rules centrally managed on Panorama and applied by the VM-Series firewall**—The next-generation firewall rules are applied by the VM-Series firewall. These rules are centrally defined and managed on Panorama using template stacks and device groups and pushed to the VM-Series firewalls. The VM-Series firewall then enforces security policy by matching on source or destination IP address—the use of dynamic address groups allows the firewall to populate the members of the groups in real time—and forwards the traffic to the filters on the NSX-V Firewall.

  To understand how the NSX-V Manager and Panorama stay synchronized with the changes in the SDDC and ensure that the VM-Series firewall consistently enforces policy, see Policy Enforcement using Dynamic Address Groups.

Policy Enforcement using Dynamic Address Groups

Unlike the other versions of the VM-Series firewall, because both virtual wire interfaces (and subinterfaces) belong to the same zone, the VM-Series firewall for NSX-V uses dynamic address groups as the traffic segmentation mechanism. A security policy rule on the VM-Series firewall for NSX-V must have the same source and destination zone, therefore to implement different treatment of traffic, you use dynamic address groups as source or destination objects in security policy rules.

 ©2024 Palo Alto Networks, Inc.

Dynamic address groups offer a way to automate the process of referencing source and/or destination addresses within security policies because IP addresses are constantly changing in a data center environment. Unlike static address objects that must be manually updated in configuration and committed whenever there is an address change (addition, deletion, or move), dynamic address groups automatically adapt to changes.

Any dynamic address groups created in a device group belonging to NSX-V configuration and configured with the match criterion **_nsx_<*dynamic address group name*>** trigger the creation on corresponding security groups on the NSX-V Manager. In an ESXi cluster with multiple customers or tenants, the ability to filter security groups for a service profile (zone on Panorama) on the NSX-V Manager allows you to enforce policy when you have overlapping IP addresses across different security groups in your virtual environment.

If, for example, you have a multi-tier architecture for web applications, on Panorama you create three dynamic address groups for the WebFrontEnd servers, Application servers and the Database servers. When you commit these changes on Panorama, it triggers the creation of three corresponding security groups on NSX-V Manager.

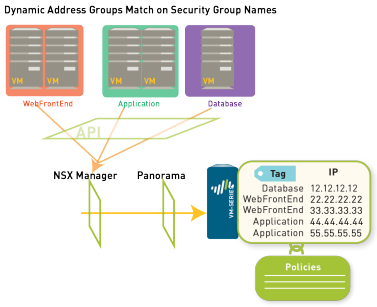

On NSX-V Manager, you can then add individual guest VMs or IP sets (IP ranges or subnets) to the appropriate security groups. Then, in security policy you can use the dynamic address groups as source or destination objects, define the applications that are permitted to traverse these servers, and push the rules to the VM-Series firewalls.

Each time a guest is added or modified in the ESXi cluster or a security group is updated or created, the NSX-V Manager uses the PAN-OS REST-based XML API to update Panorama with the IP address, and the security group to which the guest belongs. To trace the flow of information, see Dynamic Address Groups—Information Relay from NSX Manager to Panorama.

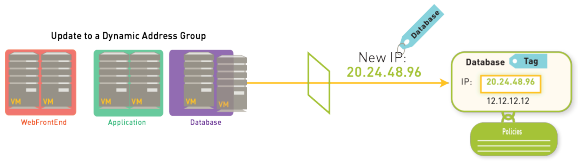

To ensure that the name of each security group is unique, the vCenter server assigns a Managed Object Reference (MOB) ID to the name you define for the security group. The syntax used to display the name of a security group on Panorama is serviceprofileid-specified_name-securitygroup-number; for example, serviceprofile13-WebFrontEnd-securitygroup-47.

When Panorama receives the API notification, it verifies/updates the IP address of each guest and the security group and the service profile to which that guest belongs. Then, Panorama pushes

 ©2024 Palo Alto Networks, Inc.

these real-time updates to all the firewalls that are included in the device group and notifies device groups in the service manager configuration on Panorama.

On each firewall, all policy rules that reference these dynamic address groups are updated at runtime. Because the firewall matches on the security group tag to determine the members of a dynamic address group, you do not need to modify or update the policy when you make changes in the virtual environment. The firewall matches the tags to find the current members of each dynamic address group and applies the security policy to the source/destination IP address that are included in the group.

What are the Benefits of the NSX-V VM-Series firewall for NSX-V Solution?

The VM-Series firewall for VMware NSX-V is focused on securing east-west communication in the software-defined data center. Deploying the firewall has the following benefits:

- **Sturdier Centralized Management**—The firewalls deployed using this solution are licensed and managed by Panorama, the Palo Alto Networks central management tool. Panorama serves as a single point of configuration for integration with NSX-V. It gives the NSX-V Manager the information is it needs to steer redirect traffic to the VM-Series firewall for inspection and enforcement. Using Panorama to manage both the perimeter and data center firewalls (the hardware-based and virtual firewalls) allows you to centralize policy management and maintain agility and consistency in policy enforcement throughout the network.
- **Automated Deployment**—The NSX-V Manager automates the process of delivering next-generation firewall security services and the VM-Series firewall allows for transparent security enforcement. When a new ESXi host is added to a cluster, a new VM-Series firewall is automatically deployed, provisioned and available for immediate policy enforcement without any manual intervention. The automated workflow allows you to keep pace with the virtual machine deployments in your data center. The hypervisor mode on the firewall removes the need to reconfigure the ports/ vswitches/ network topology; because each ESXi host has an instance of the firewall, the traffic does not need to traverse the network or be backhauled for inspection and consistent enforcement of policies.
- **Ease in Administering Tenants in Shared and Dedicated Compute Infrastructure** —This integration provides the flexibility in configuring the firewall to handle multiple zones for traffic segmentation, defining shared or specific policy sets for each tenant or sub-tenant, and includes support for overlapping IP addresses across tenants or sub-tenants. Whether you have a shared cluster and need to define tenant specific policies and logically isolate traffic for each tenant (or sub-tenant), or you have a dedicated cluster for each tenant, this solution enables you to configure the firewall for your needs. And if you need a dedicated instance of the VM-Series firewall for each tenant in a cluster that hosts the workloads for multiple tenants, you can deploy multiple instances of the VM-Series firewall on each host in an ESXi cluster. For more information, see What is Multi-Tenant Support on the VM-Series Firewall for NSX-V?
- **Tighter Integration Between Virtual Environment and Security Enforcement for Dynamic Security**—Dynamic address groups maintain awareness of changes in the virtual machines/ applications and ensure that security policy stays in tandem with the changes in the network. This awareness provides visibility and protection of applications in an agile environment.

In summary, this solution ensures that the dynamic nature of the virtual network is secured with minimal administrative overhead. You can successfully deploy applications with greater speed, efficiency, and security.

 ©2024 Palo Alto Networks, Inc.

What is Multi-Tenant Support on the VM-Series Firewall for NSX-V?

Multi-tenancy on the VM-Series firewall enables you to secure more than one *tenant* or more than one *sub-tenant*. A tenant is a customer or an organization such as Palo Alto Networks. A sub-tenant is a department or business unit within the organization such as Marketing, Accounting, or Human Resources. To allow you to secure multiple tenants, Panorama provides the flexibility to create multiple sets of security policy rules for each tenant, and multiple zones to isolate traffic from each sub-tenant and redirect traffic to the appropriately configured VM-Series firewall. You can also deploy more than one instance of the VM-Series firewall on each host within an ESXi cluster.

Panorama and managed VM-Series firewalls must be running PAN-OS 7.1 or greater to support multi-tenancy.

To deploy a multi-tenant solution, create one or more *service definition(s)* and *service profile zone(s)* on Panorama. A service definition on Panorama specifies the configuration of the VM-Series firewall using one device group and one template stack. This means that each instance of the VM-Series firewalls that is deployed using a service definition has one common set of policy rules for securing the tenants and sub-tenants in the ESXi cluster.

A service profile zone within a Panorama template stack is used to segment traffic from each sub-tenant using virtual wire subinterfaces. When you create a new service profile zone, Panorama pushes the zone as a part of the template stack configuration to the firewall, and the firewall automatically creates a pair of virtual wire subinterfaces, for example ethernet1/1.3 and ethernet 1/2.3 so that the firewall can isolate traffic for a sub-tenant. Because a template stack supports up to 32 subinterface pairs, you can logically isolate traffic and secure up to 32 sub-tenants.

Panorama registers each service definition as a service definition on the NSX-V Manager and each service profile zone as a service profile within the corresponding service definition. And, when you deploy the service definition from the NSX-V Manager, an instance of the VM-Series firewall is deployed on each host in the ESXi cluster. And you can use the steering rules defined on Panorama and applied to the NSX-V Manager to specify what traffic to redirect to the VM-Series firewall based on NSX-V security groups, and to which tenant or sub-tenant based on the service profile.

Based on your requirements, you can choose from the following multi-tenancy options:

- **Shared cluster with shared VM-Series firewalls**- Multiple tenants share the cluster and the VM-Series firewall. A single instance of the VM-Series firewall is deployed on each host in the cluster. In order to separate traffic from each tenant, you create a zone for each tenant, and you define a single, common set of policy rules to secure the virtual machines for all tenants. See Use Case: Shared Compute Infrastructure and Shared Security Policies.
- **Dedicated cluster with dedicated VM-Series firewalls**- A single tenant occupies the cluster, and a single instance of the VM-Series firewall is deployed on each host in the cluster. In this deployment, the tenant can have a single zone and a single policy set, or the tenant can have multiple zones for sub-tenants that require traffic separation (one zone per sub-tenant) and a single policy set with zone-based rules to secure traffic for each sub-tenant. Use Case: Shared Security Policies on Dedicated Compute Infrastructure.
- **Shared cluster with dedicated VM-Series firewalls**- Multiple tenants share the cluster and multiple instances of the VM-Series firewalls are deployed on each host in a cluster so that each tenant can have a dedicated instance of the VM-Series firewall. This deployment provides

 ©2024 Palo Alto Networks, Inc.

scalability and better performance on shared infrastructure for each tenant. Based on each tenant's needs, you will define two or more service definitions for the cluster.

When deploying multiple instances of the VM-Series firewall, you must ensure that each ESXi host has the sufficient CPU, memory and hard disk resources required to support the VM-Series firewalls and the other virtual machines that will be running on it.

VM-Series Firewall for NSX-V Deployment Checklist

To deploy the VM-Series firewall for NSX-V, use the following workflow:

- **Step 1: Set up the Components**—To deploy the VM-Series firewall for NSX-V, set up the following components (see What are the Components of the VM-Series for NSX-V Solution?):
  - Set up the vCenter server, install and register the NSX-V Manager with the vCenter server.

    If you have not already set up the virtual switch(es) and grouped the ESXi hosts in to clusters, refer to the VMware documentation for instructions on setting up the vSphere environment. This document does not take you through the process of setting up the VMware components of this solution.

    Unless you Enable Large Receive Offload, do not modify the default value (1500 bytes) of the MTU on the virtual Distributed Switch (vDS) in the vSphere infrastructure. Modifying the MTU to any other value causes the VM-Series firewall for NSX-V to discard packets.
  - Upgrade Panorama. If you are new to Panorama, refer to the Panorama documentation for instructions on setting up and upgrading Panorama. See Migrate Operations-Centric Configuration to Security-Centric Configuration if you choose to migrate your Operations-Centric configuration to a Security-Centric configuration format.
  - Configure an SSL/TLS Service Profile. If you are running NSX-V Manager 6.2.3 or earlier, you must configure an SSL/TLS Service profile that allows TLSv1.0 and apply it to the Panorama management interface. If you are running NSX-V Manager 6.2.4 or later, an SSL/TLS Service profile is not required.
  - Install the VMware NSX Plugin.
  - Install a License Deactivation API Key. Deleting the Palo Alto Networks Service Deployment on NSX-V Manager automatically triggers license deactivation. A license deactivation API key is required to successfully deactivate the VM-Series license.
  - Download and save the ovf template for the VM-Series firewall for NSX-V on a web server. The ovf template must match your VM-Series model. If you are using the VM-200, select the VM-100 ovf. If using the VM-1000-HV, select the VM-300 ovf.

    The NSX-V Manager must have network access to this web server so that it can deploy the VM-Series firewall as needed. You cannot host the ovf template on Panorama.

    *Give the ova filename a generic name that does not include a version number. Using a generic naming convention, such as **https://acme.com/software/PA-VM-NSX.ova** allows you to overwrite the ova each time a newer version becomes available.*
  - Register the capacity auth-code for the VM-Series firewall for NSX-V with your support account on the Support Portal. For details, see Upgrade the VM-Series Firewall.

 ©2024 Palo Alto Networks, Inc.

- **Step 2: Register**—Configure Panorama to Register the VM-Series Firewall as a Service on the NSX-V Manager. When registered, the VM-Series firewall is added to the list of network services that can be transparently deployed as a service by the NSX-V Manager. The connection between Panorama and the NSX-V Manager is also required for licensing and configuring the firewall.
  - (On Panorama) Create a service manager to enable communication between Panorama and NSX-V Manager.
  - (On Panorama) Create the service definition. If you upgrade from an earlier version, your existing service definition is automatically migrated for you.
- **Step 3: Deploy the VM-Series Firewall**—Before you can deploy the VM-Series firewall in NSX-V, each host in the cluster must have the necessary NSX-V components required to deploy the firewall.
  - (On NSX-V Manager) Define the IP address pool. An IP address from the defined range is assigned to the management interface of each instance of the VM-Series firewall.

    The NSX-V Manager uses the IP address as a match criterion to steer traffic to the VM-Series firewall. If VMware tools is not installed on the guest, see Steer Traffic from Guests that are not Running VMware Tools. This is not required if you are running NSX-V Manager 6.2.4 or later.
  - (On NSX-V Manager) Prepare the ESXi host for the VM-Series firewall.
  - (On NSX-V Manager) Deploy the VM-Series firewall. The NSX-V Manager automatically deploys an instance of the VM-Series firewall on each ESXi host in the cluster.
  - (On NSX-V Manager) Add VMs to the relevant security groups.
  - (On Panorama) Apply policies to the VM-Series firewall. From Panorama, you define, push, and administer policies centrally on all the VM-Series firewalls. This centralized administration mechanism allows you to secure guests/applications with minimal administrative intervention.
- **Step 4: Create Security Groups and Steering Rules**—How you choose to deploy the security groups and steering rules depends on whether your deployment focus is Security Centric or Operations Centric.

  In a Security Centric deployment, your security administrator creates the security group and steering rules in Panorama. You might start with an existing set of security policies and a set of named source and destination groups. Any new dynamically deployed applications fit into predefined security policies defined on Panorama. Panorama pushes these named groups to NSX-V Manager, where the virtualization administrator picks up the group names and defines which VMs go into them.

  In an Operations Centric deployment, security groups are defined by a virtualization administrator based upon the need to classify and categorize VM workloads. In this case, security groups are defined and populated in the NSX-V Manager. Security groups created in NSX-V Manager must be associated with dynamic address groups on Panorama, which is

 ©2024 Palo Alto Networks, Inc.

completed after the firewalls are deployed. In this case, NSX-V base functionality is deployed first and the VM-Series firewalls are added later.

You must decide whether a Security Centric or an Operations Centric deployment is right for your NSX-V environment before continuing. This document describes the procedure for a Security Centric deployment.

Security Centric—Create the service definition(s) that specify the configuration for the VM-Series firewall, create dynamic address groups, and create policies to redirect traffic to the VM-Series firewall. See Create Security Groups and Steering Rules in a Security Centric Deployment.

- (On Panorama) Set up the dynamic address groups that map to security groups on NSX-V Manager. A security group assembles the specified guests/applications so that you can apply policy to the group.
- (On Panorama) Create the security policy rules to redirect traffic to the Palo Alto Networks service profile.

Operations Centric—On the NSX-V Manager, create security groups and policies to redirect traffic to the VM-Series firewall. See Create Security Groups and Steering Rules in an Operations Centric Deployment.

- (On NSX-V Manager) Set up the security groups. A security group assembles the specified guests/applications so that you can apply policy to the group.
- (On NSX-V Manager) Create the NSX-V Firewall policies to redirect traffic to the Palo Alto Networks service profile.

- **Step 5: Monitor and Maintain Network Security**—Panorama provides a comprehensive, graphical view of network traffic. Using the visibility tools on Panorama—the Application Command Center (ACC), logs, and the report generation capabilities—you can centrally analyze, investigate and report on all network activity, identify areas with potential security impact, and translate them into secure application enablement policies. Refer to the Panorama Administrator's Guide for more information.

The following additional tasks are not required parts of the main VM-Series for NSX-V deployment procedure and should only be completed if and when necessary for your deployment.

- **Upgrade the Software Version—**When upgrading the VM-Series firewalls for NSX-V, you must first upgrade Panorama before upgrading the firewalls. To upgrade the firewalls, see Upgrade the PAN-OS Software Version (VM-Series for NSX).

  - *For upgrading the PAN-OS version on the firewall, do not modify the **VM-Series OVA URL** in **Panorama** > **VMware Service Manager**.*
  - *Do not use the VMware snapshots functionality on the VM-Series firewall for NSX-V. Snapshots can impact performance and result in intermittent and inconsistent packet loss. See VMware's best practice recommendation with using snapshots. If you need configuration backups, use Panorama or **Export named configuration snapshot** from the firewall (**Device** > **Set up** > **Operations**). Using the Export named configuration snapshot exports the active configuration (running-config.xml) on the firewall and allows you to save it to any network location.*

- Migrate from Operations-Centric configuration to Security-Centric configuration—If you upgrade your existing Operations-Centric VM-Series firewall for NSX-V deployment and plan

 ©2024 Palo Alto Networks, Inc.

to use the Security Centric workflow going forward, Migrate Operations-Centric Configuration to Security-Centric Configuration.

If you need to reinstall or remove the VM-Series from your NSX-V deployment, see the How to Remove VM-Series Integration from VMware NSX-V knowledge base article.

Install the VMware NSX Plugin

To deploy the VM-Series for NSX solution, you must install the VMware NSX plugin on Panorama.

*If another version of the plugin is currently installed, selecting **Install** removes it and installs the selected version.*

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Download the plugin.

1. Select **Panorama** > **Plugins**.
2. Select **Check Now** to retrieve a list of available updates.
3. Select **Download** in the Action column to download the plugin.
4. Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

When installing the plugin on Panoramas in an HA pair, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | If you are upgrading your version of the VMware NSX plugin, complete a manual configuration sync.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Managers**.
2. Select **NSX Config-Sync** in the Action column.
3. Click **Yes**.
4. When the sync is complete, click **OK**.

Register the VM-Series Firewall as a Service on the NSX-V Manager

You need to enable communication between Panorama and the NSX-V Manager and then register the VM-Series firewall as a service on the NSX-V Manager. When registered, the VM-Series firewall is added to the list of network services that can be transparently deployed as a service by the NSX-V Manager.

- Enable Communication Between the NSX-V Manager and Panorama
- Create Template(s) and Device Group(s) on Panorama
- Create the Service Definitions on Panorama

Enable Communication Between the NSX-V Manager and Panorama

To automate the provisioning of the VM-Series firewall for NSX-V, enable communication between the NSX-V Manager and Panorama. This is a one-time setup, and only needs to be modified if the IP address of the NSX-V Manager changes or if the capacity license for deploying the VM-Series firewall is exceeded.

STEP 1 | (Optional) Bypass proxy server settings, configured on Panorama under **Panorama** > **Setup** > **Services** > **Proxy Server**, for communication between Panorama and NSX-V Manager. This command allows Panorama to communicate directly with NSX-V Manager while maintaining proxied communication for other services. This feature requires Panorama plugin for VMware NSX 2.0.5.

1. Log in to the Panorama CLI.
2. Execute the following command to enable or disable proxy bypass.

   ```
   admin@Panorama> request plugins vmware_nsx global proxy bypass
   {yes | no}
   ```

   Select **yes** to enable proxy bypass and **no** to disable proxy bypass.

STEP 2 | Log in to the Panorama web interface.

Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<*IP address*>).

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Set up access to the NSX-V Manager.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Managers** and click **Add**.
2. Enter the **Service Manager Name**.

   On the NSX-V Manager, this name displays in the Service Manager column on **Networking & Security** > **Service Definitions** > **Service Managers**.
3. (Optional) Add a **Description** that identifies the VM-Series firewall as a service.
4. Enter the **NSX Manager URL**—IP address or FQDN—at which to access the NSX-V Manager.
5. Enter the **NSX Manager Login** credentials—the username and password for your Enterprise Administrator role on NSX Manager. This allows Panorama to authenticate with the NSX-V Manager.

   The ampersand (&) special character is not supported in the NSX-V manager account password. If a password includes an ampersand, the connection between Panorama and NSX-V manager fails.

   If you change your NSX-V Manager login password, ensure that you update the password on Panorama immediately. An incorrect password breaks the connection between Panorama and NSX-V Manager. Panorama does not receive updates about changes to your deployment while disconnected from NSX-V Manager.
6. Click **OK**.

STEP 4 | Commit your changes to Panorama.

Select **Commit** and Commit Type: **Panorama**.

STEP 5 | Verify the connection status on Panorama.

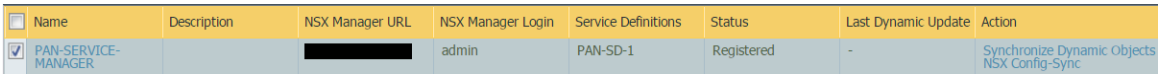

To view the connection status between Panorama and the NSX-V Manager.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Managers**.
2. Verify the message in the **Status** column.

   When the connection is successful, the status displays as **Registered**. This indicates that Panorama and the NSX-V Manager are in sync and the VM-Series firewall is registered as a service on the NSX-V Manager.

   The unsuccessful status messages are:

   - **Not connected**: Unable to reach/establish a network connection to the NSX-V Manager.
   - Invalid Credentials: The access credentials (username and/or password) are incorrect.
   - **Out of sync**: The configuration settings defined on Panorama are different from what is defined on the NSX-V Manager.Click the link for details on the reasons for failure. For example, NSX-V Manager may have a service definition with the same name as defined on Panorama. To fix the error, use the service definition name listed in the

 ©2024 Palo Alto Networks, Inc.

error message to validate the service definition on the NSX-V Manager. Until the configuration on Panorama and the NSX-V Manager is synchronized, you cannot add a new service definition on Panorama.

- **No service/ No service profile**: Indicates an incomplete configuration on the NSX-V Manager.

If you make a change and need to manually sync, see 10

STEP 6 | Verify that the firewall is registered as a service on the NSX-V Manager.

1. On the vSphere web client, select **Networking & Security** > **Service Definitions** > **Service Managers**.

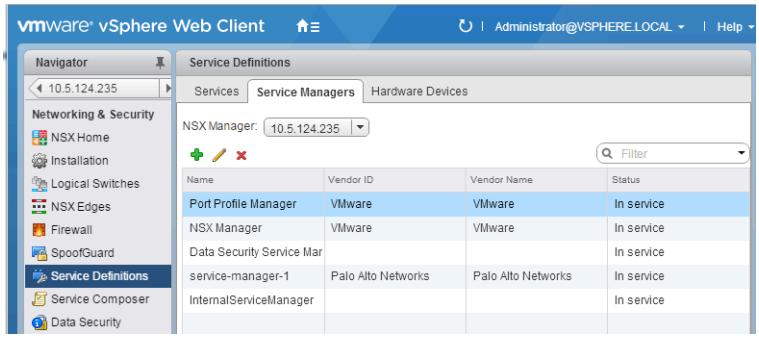

2. Verify that **Palo Alto Networks** displays as a vendor in the list of services available for installation.

STEP 7 | If you are running VMware NSX plugin 2.0.4 or later, you can configure Panorama to automatically synchronize dynamic objects with NSX-V manager as if you issued an **Synchronize Dynamic Objects**. By default, the DAG Sync interval is disabled and the value is set to zero (0). To enable the DAG Sync, set the interval between one hour and 72 hours. Setting a value of zero hours disables the DAG sync. To configure or disable the interval, complete the following procedure.

1. Log in to the Panorama CLI.
2. Execute the following command.

   ```
   request plugins vmware_nsx nsx_v dag-sync-interval interval <interval-in-hours>
   ```

   You can view the configured value with the following show command.

   ```
   show plugins vmware_nsx nsx_v dag-sync-interval
   ```

STEP 8 | (Optional) In large NSX-V environments with tens of thousands of IP addresses, allowing Panorama enough time to retrieve IP address updates from NSX-V Manager is essential. You can now configure the amount of time—up to 10 minutes—Panorama has to retrieve updates from NSX-V Manager. By default, Panorama waits up to two minutes (120 seconds) to get

 ©2024 Palo Alto Networks, Inc.

IP address updates from NSX-V Manager. However, if Panorama does not retrieve all the IP address updates within the alloted two minutes, Panorama times out and the update fails.

You can determine if you are experiencing curl call failures through the Panorama error log. Curl call failures return the following message.

```
2019-05-23 06:50:15.780 -0700 ERROR: Curl call to NSX Manager
 failed
```

Complete the following procedure to increase the amount of time Panorama has to process updates. This feature requires Panorama plugin for VMware NSX 2.0.5.

1. Log in to the Panorama CLI.
2. Execute the following command to set the curl call timeout. You can set the time out from 30 seconds to 600 seconds (10 minutes).

   admin@Panorama> request plugins vmware_nsx global curl-timeout timeout <seconds>

   If Panorama is part of an HA pair, configure the same timeout value on the active and passive Panorama peers.

Create Template(s), Template Stack(s), and Device Group(s) on Panorama

To manage the VM-Series firewalls for NSX-V using Panorama, the firewalls must belong to a device group and a template that is a member of a template stack. Device groups allow you to assemble firewalls that need similar policies and objects as a logical unit; the configuration is defined using the **Objects** and **Policies** tabs on Panorama. Use template stacks to configure the settings that are required for the VM-Series firewalls to operate on the network and associate; the configuration is defined using the **Device** and **Network** tabs on Panorama. And each template stack with zones used in your NSX-V configuration on Panorama must be associated with a service definition; at a minimum, you must create a zone within the template stack so that the NSX-V Manager can redirect traffic to the VM-Series firewall.

Each virtual wire zone belonging to the NSX-related template becomes available as a *service profile* on the Service Composer on the NSX-V Manager. When you create NSX-related zone on Panorama, Panorama pushes the zone as a part of the template stack configuration to the firewall, and the firewall automatically creates a pair of virtual wire subinterfaces, for example ethernet1/1.3 and ethernet 1/2.3, to isolate traffic for a tenant or sub-tenant. On the firewall, you can then Create Security Groups and Steering Rules to secure traffic that arrives on the virtual wire subinterface pair that maps to the zone.

If you are new to Panorama, refer to the Panorama Administrator's Guide for instructions on setting up Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Add a device group or a device group hierarchy.

1. Select **Panorama** > **Device Groups**, and click **Add**. You can also create a device group hierarchy.
2. Enter a unique **Name** and a **Description** to identify the device group.
3. Click **OK**.

   After the firewalls are deployed and provisioned, they will display under **Panorama** > **Managed Devices** and will be listed in the device group.
4. Click **Commit** and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 2 | Add a template.

1. Select **Panorama** > **Templates**, and click **Add**.
2. Enter a unique **Name** and a **Description** to identify the template.
3. Click **OK**.
4. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 3 | Add a template stack.

1. Select **Panorama** > **Templates**, and click **Add Stack**.
2. Enter a unique **Name** and a **Description** to identify the template stack.
3. Click **Add** under Templates and select the template you created above.
4. Click **OK**.
5. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 4 | Create the zone(s) for each template.

Each zone is mapped to a service profile on NSX-V Manager. To qualify, a zone must be of the virtual wire type and a template associated with a service definition.

For a single-tenant deployment, create one zone. If you have multi-tenant deployment, create a zone for each sub-tenant.

You can add up to 32 zones in each template.

1. Select **Network** > **Zones**.
2. Select the correct template in the **Template** drop-down.
3. Select **Add** and enter a zone **Name**.
4. Set the interface **Type** to **Virtual Wire**.
5. Click **OK**.
6. Verify that the zones are attached to the correct template.
7. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

   Panorama creates a corresponding service profile on NSX-V Manager for each qualified zone upon commit.

 ©2024 Palo Alto Networks, Inc.

Create the Service Definitions on Panorama

A service definition specifies the configuration for the VM-Series firewalls installed on each host in an ESXi cluster. The service definition must include the device group, the license auth-codes for deploying the VM-Series firewalls, and a template stack with one or more NSX-V service profile zones. Typically, you create a service definition for the VM-Series firewall in an ESXi cluster. If you have different ESXi clusters that have workloads that require the VM-Series firewall to handle traffic differently, you can create multiple service definitions on Panorama.

On a Panorama commit, each service definition is registered on the NSX-V Manager. On registration with the NSX-V Manager, the NetX API implementation makes each zone (defined within the template stack) available for redirecting traffic. When you deploy the VM-Series firewalls, you can select the profile name for the VM-Series firewall(s) to which you want to redirect traffic from the objects in NSX-V security groups. The appropriately configured firewall can then inspect the traffic and enforce policy from the virtual machines that belong to the NSX-V security groups.

STEP 1 | (Optional) Configure a Notify Group

Create a notify group by specifying devices groups that should be notified of changes in the virtual environment. The firewalls included in the specified device groups receive a real-time update of security groups and IP addresses of guest VMs in them. The firewalls use this update to determine the most current list of members that constitute dynamic address groups referenced in policy

1. Select **Panorama** > **VMware** > **NSX-V** > **Notify Group** and click **Add**.
2. Give your Notify Group a descriptive **Name**.
3. Select the boxes of all devices groups that should be notified of changes to the virtual environment. If a device group does not have a check box available, it means that the device group is automatically included by virtue of the device group hierarchy.
4. Click **OK**.

STEP 2 | Add a new service definition.

You can create up to 32 service definitions on Panorama.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Definitions**.
2. Select **Add** to create a new service definition. The maximum number of characters in a service definition name is 40.

   On the NSX-V Manager, this service definition name displays in the Services column on **Networking & Security** > **Service Definitions** > **Services**.
3. (Optional) Add a **Description** that identifies the function or purpose for the VM-Series firewalls that will be deployed using this service definition.

STEP 3 | Assign a device group and a template stack to the service definition.

Make sure to Create the zone(s) for each template stack.

Because the firewalls deployed in this solution will be centrally administered from Panorama, you must specify the **Device Group** and the **Template Stack** that the firewalls belong to. All the

 ©2024 Palo Alto Networks, Inc.

firewalls that are deployed using this service definition belong to the specified template stack and device group.

1. Select the device group or device group hierarchy in the **Device Group** drop-down.
2. Select the template stack in the **Template** drop-down.

   You cannot reuse a template stack or a device group assigned to one service definition in another service definition.

STEP 4 | Specify the location of the OVF file.

Download the zip file, unzip it to extract and save the .ovf, mf and .vmdk files to the same directory. Both the files are used to deploy each instance of the firewall.

If needed, modify the security settings on the server so that you can download the file types. For example, on the IIS server modify the Mime Types configuration; on an Apache server edit the .htaccess file.

Do not change the Panorama service definition OVF path after a successful NSX Service Deployment of VM-Series firewalls. Changing the OVF path, after a successful VM-Series firewall deployment, can result in a NSX Service Deployment failed state. You may resolve this failure in NSX-V Manager, however this may cause all VM-Series firewalls to redeploy.

It is recommended that you use an OVF path name that scales and allows you to change the base image without impacting your deployed firewalls. Instead of a path such as **https://acme.com/software/PA-VM-NST.9.1.0.ovf**, use something such as **https://acme.com/software/PanoSvcDef1-Cluster1.ovf**. Using a static path reference will eliminate any future need to change the OVF path. It is recommended to create a path for each Panorama service definition (vSphere cluster) in your deployment and change the PAN-OS base images references on the web server as needed.

To use a different OVF for an existing Service Definition, change the reference link on the web server to point to another OVF. Changing the web server OVF reference links will not

 ©2024 Palo Alto Networks, Inc.

impact existing NSX Service Deployments. The new OVF path will take affect on new service deployments or any hosts introduced into the vSphere Cluster.

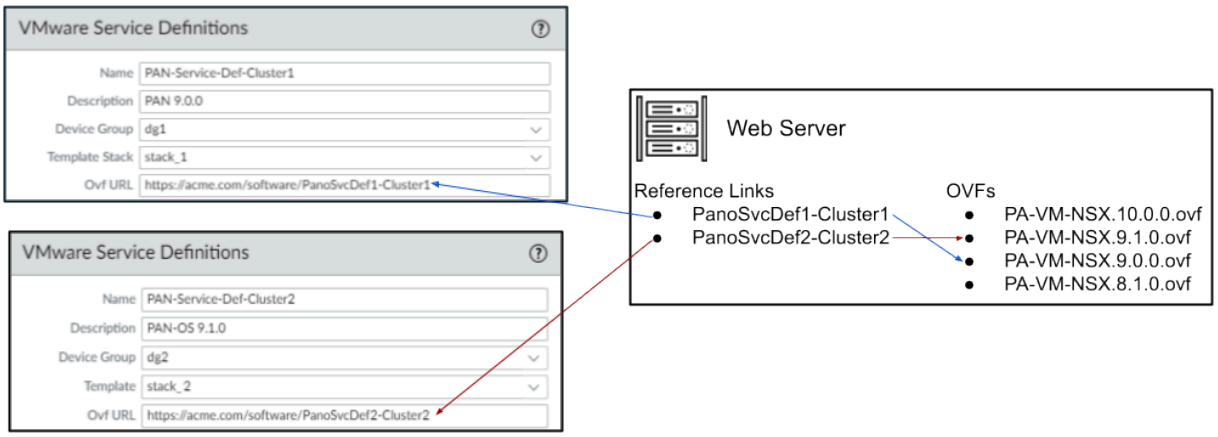

In **VM-Series OVF URL**, add the location of the web server that hosts the ovf file. Both http and https are supported protocols. For example, enter **https://acme.com/software/PA-VM-NSX.9.1.0.ovf**

Select the ovf file that matches the VM-Series model you plan to deploy. For the VM-200, use vm100.ovf. For the VM-1000-HV, use vm300.ovf.

You can use the same ovf version or different versions across service definitions. Using different ovf versions across service definitions allows you to vary the PAN-OS version on the VM-Series firewalls in different ESXi clusters.

STEP 5 | (Optional) Select a Notify Group.

To create context awareness between the virtual and security environments so that policy is consistently applied to all traffic steered to the firewalls, select the device groups to notify when there are changes in the virtual environment.

Select each device group to which you want to enable notifications in the **Notify Device Groups** drop-down. If a device group does not have a checkbox available, it means that the device group is automatically included by virtue of the device group hierarchy.

The firewalls included in the specified device groups receive a real-time update of security groups and IP addresses. The firewalls use this update to determine the most current list of members that constitute dynamic address groups referenced in policy.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | To automatically retrieve a device certificate when the VM-Series firewall is deployed by NSX Manager, configure the device certificate.

1. If you have not done so already, log in to the Customer Support Portal and generate a Registration PIN and PIN ID.
2. Under **Device Certificate**, click **Enable**.
3. Copy the PIN ID and enter it into the **Device Certificate PIN ID** field.
4. Reenter the PIN ID into the **Confirm Device Certificate PIN ID** field.
5. Copy the Registration ID and enter it into the **Device Certificate PIN Value** field.
6. Reenter the Registration ID into the **Confirm Device Certificate PIN Value** field.

STEP 7 | Save the service definition and attach it to the service manager.

1. Click **OK**.
2. Select **Panorama** > **VMware** > **NSX-V** > **Service Manager** and click the link of the service manager name.
3. Under Service Definitions, click **Add** and select your service definition from the drop-down.
4. Click **OK**.
5. Select **Commit** and Commit Type: **Panorama**.

   Committing the changes triggers the process of registering each service definition as a security service on the NSX-V Manager.

STEP 8 | Add the authorization code to license the firewalls.

The auth-code must be for the VM-Series model NSX bundle; for example, PAN-VM-300-PERP- BND-NSX.

Verify that the order quantity/ capacity is adequate to support the number of firewall you need to deploy in your network.

1. Select **Panorama** > **Device Groups** and choose the device group you associated with the service definition you just created.
2. Under Dynamically Added Device Properties, add the authorization code you received with your order fulfillment email and select a PAN-OS software version from the SW Version drop-down.

   When a new firewall is deployed under NSX-V and added to the selected device group, the authorization code is applied and the firewall is upgraded to the select version of PAN-OS.

   On the support portal, you can view the total number of firewalls that you are authorized to deploy and the ratio of the number of licenses that have been used to the total number of licenses enabled by your auth-code.
3. Synchronize the configuration between Panorama and the NSX-V Manager.
   1. Select **Panorama** > **VMware** > **NSX-V** > **Service Managers**.
   2. Select **NSX Config-Sync** under the Actions column.
   3. Click **Yes** to confirm the sync.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Verify that the service definition and the NSX-V service profile that you defined on Panorama are registered on the NSX-V Manager.

1. On the NSX-V Manager, to verify that the service definition is available, select **Networking & Security** > **Service Definitions** > **Services**. The service definition is listed as a Service on the NSX-V Manager.

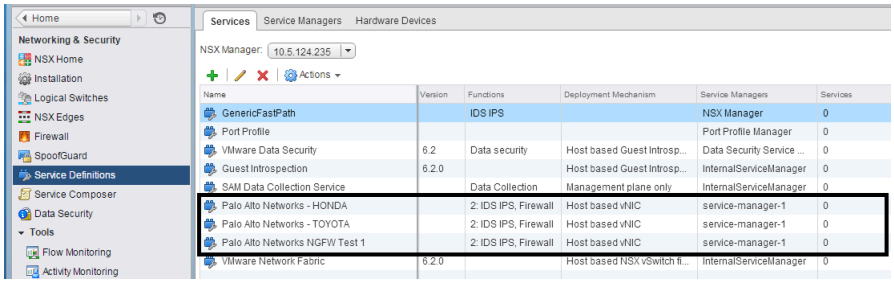

2. To verify that the zones are available on the NSX-V Manager:
   1. Select **Networking and Security** > **Service Composer** > **Security Policies**, and click **Create Security Policy**.
   2. Select **Network Introspection Services**, and click **Add**.
   3. In the **Service Name** drop-down, select a Palo Alto Networks service that you verified in the step above.
   4. In the **Profile** drop-down, verify that you can view all the zones you defined for that service definition on Panorama.

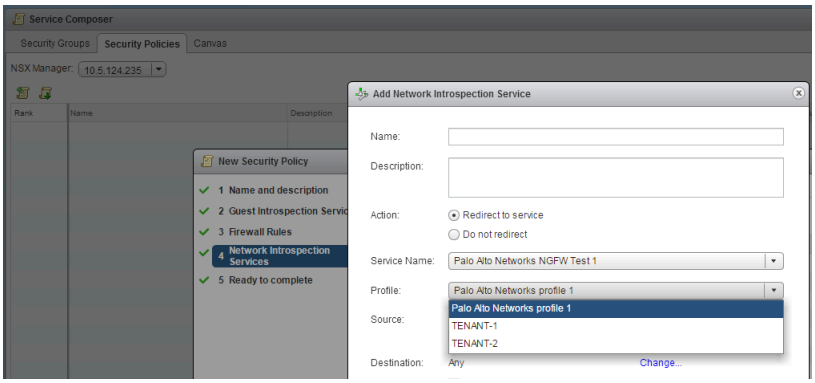

STEP 10 | (Optional) Synchronize the configuration between Panorama and the NSX-V Manager.

If you add or update the service definitions configured on Panorama, select **NSX Config Sync** in the Action column under **Panorama** > **VMware** > **NSX-V** > **Service Managers** to synchronize the changes on the NSX-V Manager.

This link is not available if you have any pending commits on Panorama.

If the synchronization fails, view the details to know whether to fix the error on Panorama or on the NSX-V Manager. For example, if you delete a service definition on Panorama, but the service definition cannot be deleted from the NSX-V Manager because it is referenced in a rule on the NSX-V Manager, the synchronization will fail with an error message that indicates the reason for failure.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall

After registering the VM-Series firewall as a service (Palo Alto Networks NGFW) on the NSX-V Manager and creating security groups and steering rules, complete the following tasks on the NSX-V Manager.

- Define an IP Address Pool (Required only if the management interface is not configured for DHCP)
- Prepare the ESXi Host for the VM-Series Firewall
- Deploy the Palo Alto Networks NGFW Service
- Enable Large Receive Offload

Support for vMotion of guest virtual machines in the vSphere/NSX-V Environment

When a guest VM is vMotioned from one host to another within a cluster, the target host NSX-V distributed firewall will steer all new sessions to the VM-Series firewall on the destination host. To ensure that all active (existing sessions) remain uninterrupted during and after the guest vMotion, the NSX-V Manager polls the VM-Series firewall for existing allowed sessions and then shares these sessions with the NSX-V distributed firewall on the destination host. All existing sessions that were allowed by the original VM-Series will be allowed by the NSX-V distributed firewall (filtering module) on the destination host without steering to the target host VM-Series firewall to prevent session loss.

The VM-Series firewall runs as a service on each host of the cluster and therefore is never vMotioned.

Define an IP Address Pool

You can configure the management interface on the VM-Series firewall to use an IP address from a static IP pool or to be a DHCP client.

If you opt to use an IP pool, which is a range of (static) IP addresses that are reserved for establishing management access to the VM-Series firewalls, when the NSX-V Manager deploys a new VM-Series firewall, the first available IP address from this range is assigned to the management interface of the firewall.

STEP 1 | In the **Networking & Security Inventory**, select the **NSX Manager**, and double click to open the configuration details of the NSX-V Manager.

STEP 2 | Select **Manage** > **Grouping Objects** > **IP Pools**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Click **Add IP Pool** and specify the network access details requested in the screen including the range of static IP addresses that you want to use for the Palo Alto Networks NGFW.

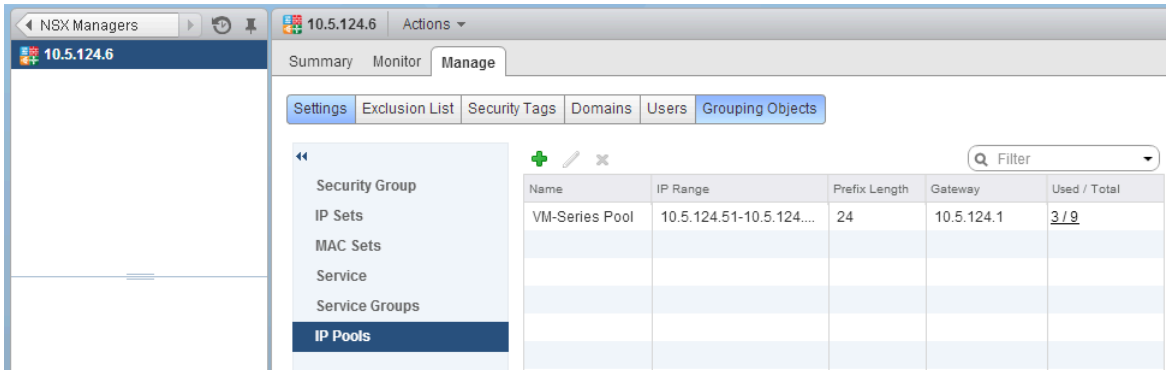

Prepare the ESXi Host for the VM-Series Firewall

Before you deploy the VM-Series firewall, each host in the cluster must have the necessary NSX-V components that allow the NSX-V firewall and the VM-Series firewall to work together. The NSX-V Manager will install the components— the Ethernet Adapter Module (.eam) and the SDK — required to deploy the VM-Series firewall.

STEP 1 | On the NSX-V Manager, select **Networking and Security** > **Installation** > **Host Preparation**.

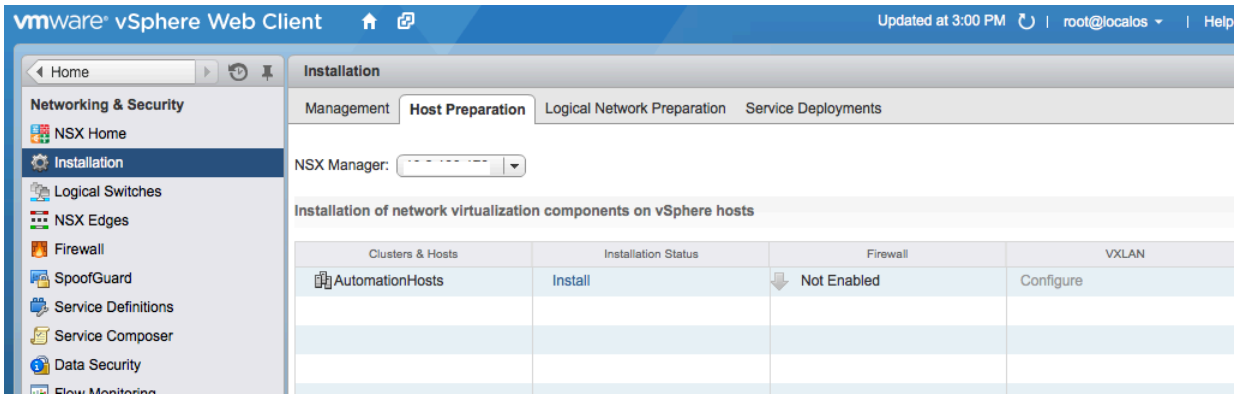

STEP 2 | Click **Install** and verify that the installation status is successful.

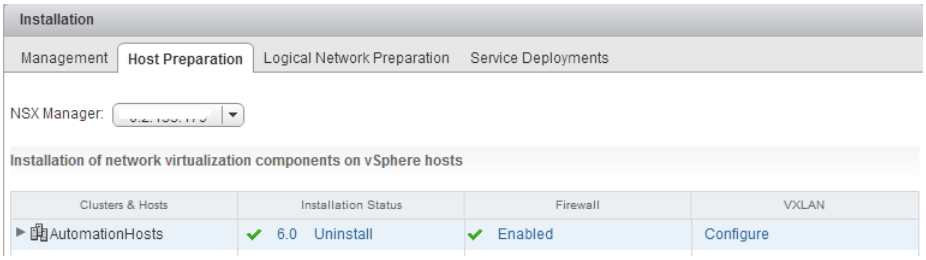

As new ESXi hosts are added to a cluster, this process is automated and the necessary NSX-V components are automatically installed on each guest on the ESXi host.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | If the Installation Status is not ready or a warning displays on screen, click the **Resolve** link. To monitor the progress of the re-installation attempt, click the **More Tasks** link and look for the successful completion of the following tasks:

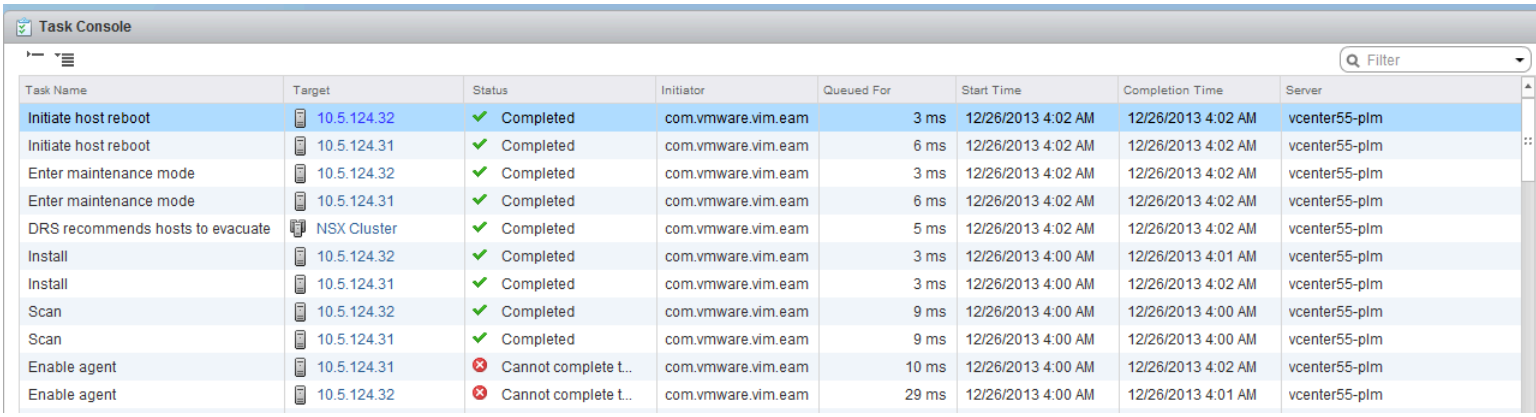

Deploy the Palo Alto Networks NGFW Service

Use the following steps to automate the process of deploying an instance of the VM-Series firewall for NSX-V on each ESXi host in the specified cluster.

STEP 1 | Select **Networking and Security** > **Installation** > **Service Deployments**.

STEP 2 | Click **New Service Deployment** (green plus icon), and select the service definition for the Palo Alto Networks next generation firewall you want to deploy, **Palo Alto Networks NGFW** service in this example. Click **Next**.

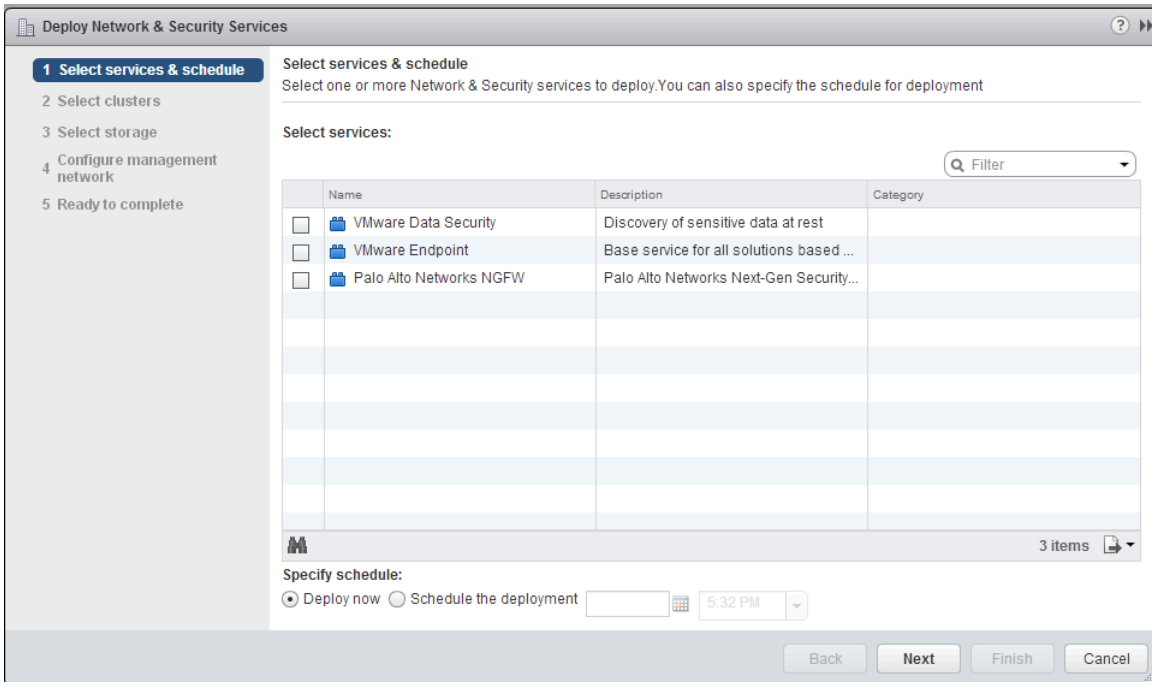

STEP 3 | Select the **Datacenter** and the cluster(s) on which the service will be deployed. One instance of the firewall will be deployed on each host in the selected cluster(s).

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Select the datastore from which to allocate disk space for the firewall. Select one of the following options depending on your deployment:

- If you have allocated shared storage for the cluster, select an available shared datastore.
- If you have not allocated shared storage for the cluster, select the **Specified-on-host** option. Be sure to select the storage on each ESXi host in the cluster. Also select the network that will be used for the management traffic on the VM-Series firewall.

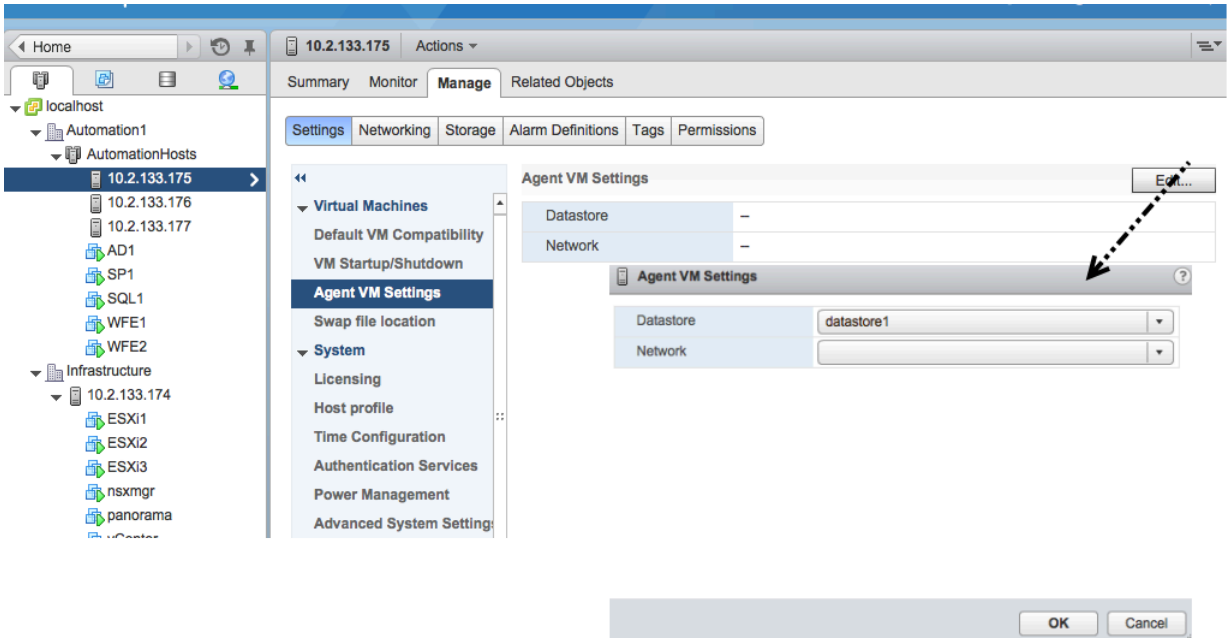

STEP 5 | Select the port group that provides management network traffic access to the firewall.

STEP 6 | Select the IP address pool assignment.

- **Use IP Pool** (Define an IP Address Pool) from which to assign a management IP address for each firewall when it is being deployed.
- **Use DHCP** on the management interface.

If you use an IP pool, on deployment, the display name for the VM-Series firewall on Panorama includes the hostname of the ESXi host. For example: PA-VM:10.5.1.120.

If you use DHCP, the display name for the VM-Series firewall does not include the name of the ESXi host.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Review the configuration and click **Finish**.

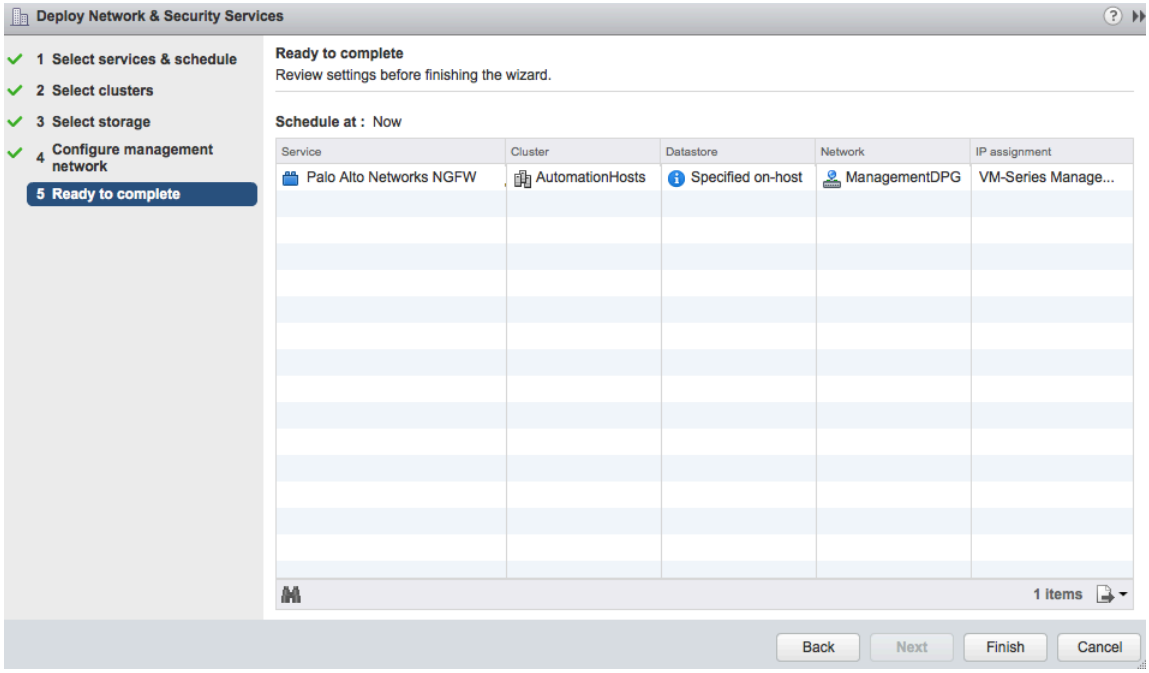

STEP 8 | Verify that the NSX-V Manager reports the **Installation Status** as **Successful**. This process can take a while; click the **More tasks** link on vCenter to monitor the progress of the installation.

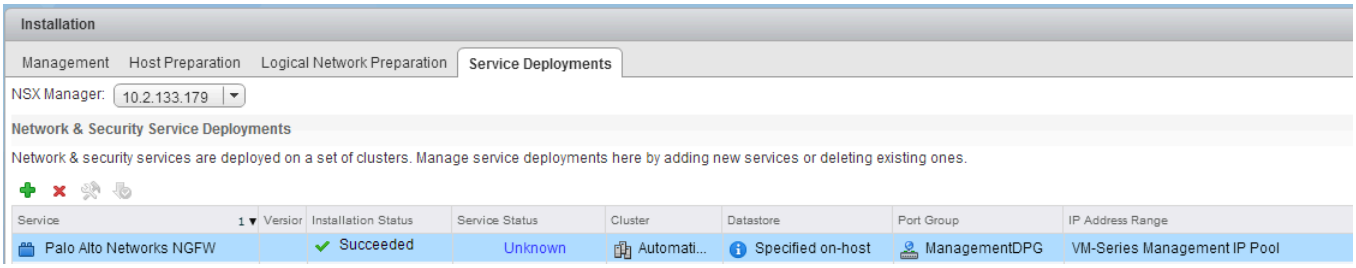

*If the installation of VM-Series fails, the error message is displayed on the Installation Status column. You can also use the **Tasks** tab and the **Log Browser** on the NSX-V Manager to view the details for the failure and refer to the VMware documentation for troubleshooting steps.*

STEP 9 | Verify that the firewall is successfully deployed.

1. On the vCenter server, select **Hosts and Clusters** to check that every host in the cluster(s) has one instance of the firewall.
2. View the management IP address(es) and the PAN-OS version running on the firewall directly from vCenter server. VMware Tools is bundled with the PAN-OS software image and is automatically enabled when you launch the VM-Series firewall.

   With VMware Tools, you can view resource utilization metrics on hard disk, memory, and CPU, and use these metrics to enable alarms or actions on the vCenter server. The heartbeats allow you to verify that the firewall is live and trigger actions to ensure high

 ©2024 Palo Alto Networks, Inc.

availability. You can also perform a graceful shutdown and restart of the firewall using the power off function on vCenter.

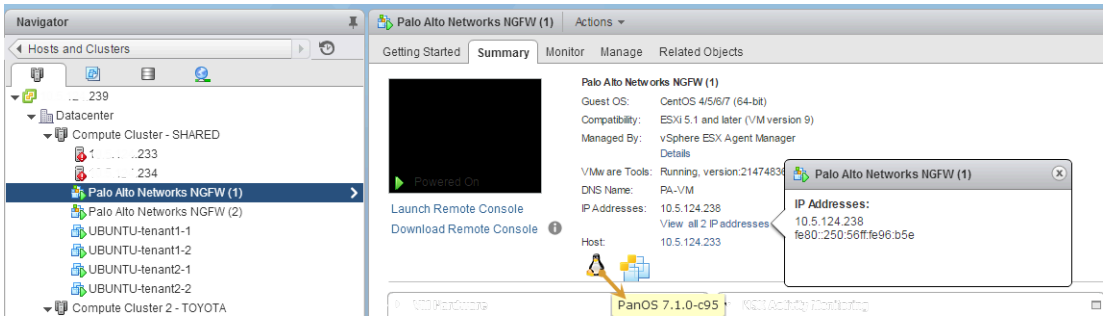

STEP 10 | Access the Panorama web interface to make sure that the VM-Series firewalls are connected and synchronized with Panorama.

1. Select **Panorama** > **Managed Devices** to verify that the firewalls are connected and synchronized.

   If the firewall gets its IP address from an IP Pool, the **Display Name** for the firewall includes the hostname of the ESXi server on which it is deployed, for example PA-VM:ESX1.Sydney. If the firewall gets a DHCP assigned IP address, the hostname of the ESXi server does not display.

   *If the ESXi server hostname is longer than 32 characters, the hostname will not be displayed in Panorama. Instead, only **PA-VM** is displayed.*

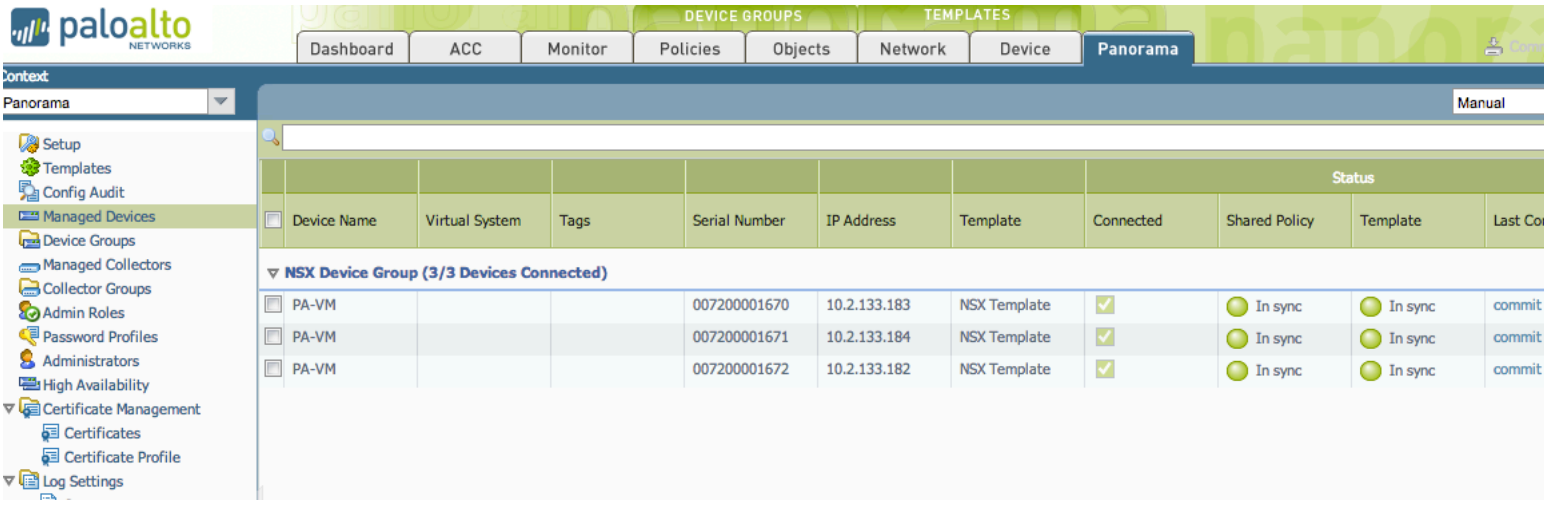

2. Click Commit, and select Commit Type as **Panorama**.

   *A periodic Panorama commit is required to ensure that Panorama saves the device serial numbers to configuration. If you reboot Panorama without committing the changes, the managed devices will not connect back to Panorama; although the Device Group will display the list of devices, the devices will not display in **Panorama** > **Managed Devices**.*

 ©2024 Palo Alto Networks, Inc.

STEP 11 | Verify that the capacity license is applied and apply any additional licenses that you have purchased. At a minimum, you must activate the support license on each firewall.

When Panorama does not have internet access (Offline), you must manually license each firewall, and then add the serial number of the firewall to Panorama so that it is registered as a managed device, and can receive the template stack and device group settings from Panorama. See Activate the License for the VM-Series Firewall for VMware NSX *for more information.*

1. Select **Panorama** > **Device Deployment** > **Licenses** to verify that the VM-Series capacity license is applied.

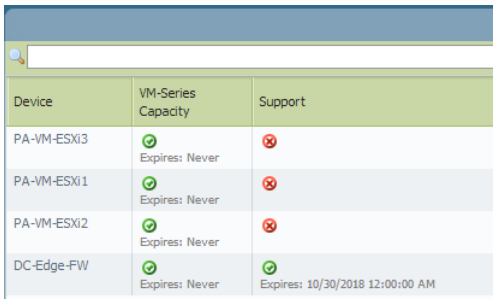

2. To apply additional licenses on the VM-Series firewalls:
   - Click **Activate** on **Panorama** > **Device Deployment** > **Licenses**.
   - Find or filter for the firewall, and in the **Auth Code** column, enter the authorization code for the license to activate. Only one authorization code can be entered at a time, for each firewall.

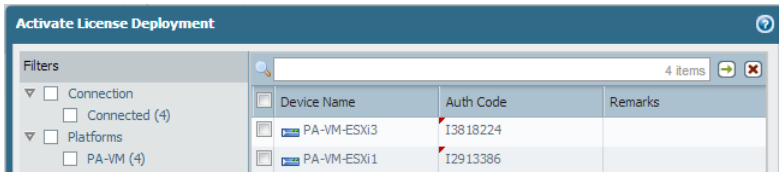

3. Click **Activate**, and verify that the result of the license activation was successful.

STEP 12 | (Optional) Upgrade the PAN-OS version on the VM-Series firewalls, see Upgrade the PAN-OS Software Version (VM-Series for NSX).

STEP 13 | Add guest VMs to the right security groups for traffic from those VMs to be redirected to the VM-Series firewall.

1. Log in to vCenter.
2. Select **Networking & Security** > **Service Composer** > **Security Groups**.
3. Highlight the security group to which you want to assign guest VMs and click the **Edit Security Group** icon.
4. Select **Define dynamic membership** and click the + icon.
5. Click **Add**.
6. Define the dynamic membership criteria that the guest VMs must meet to be part of the selected security group. The criteria you use depends on your network deployment.

 ©2024 Palo Alto Networks, Inc.

For example, you might choose to group VMs by an Entity such as Logical Switch or Distributed Port Group.

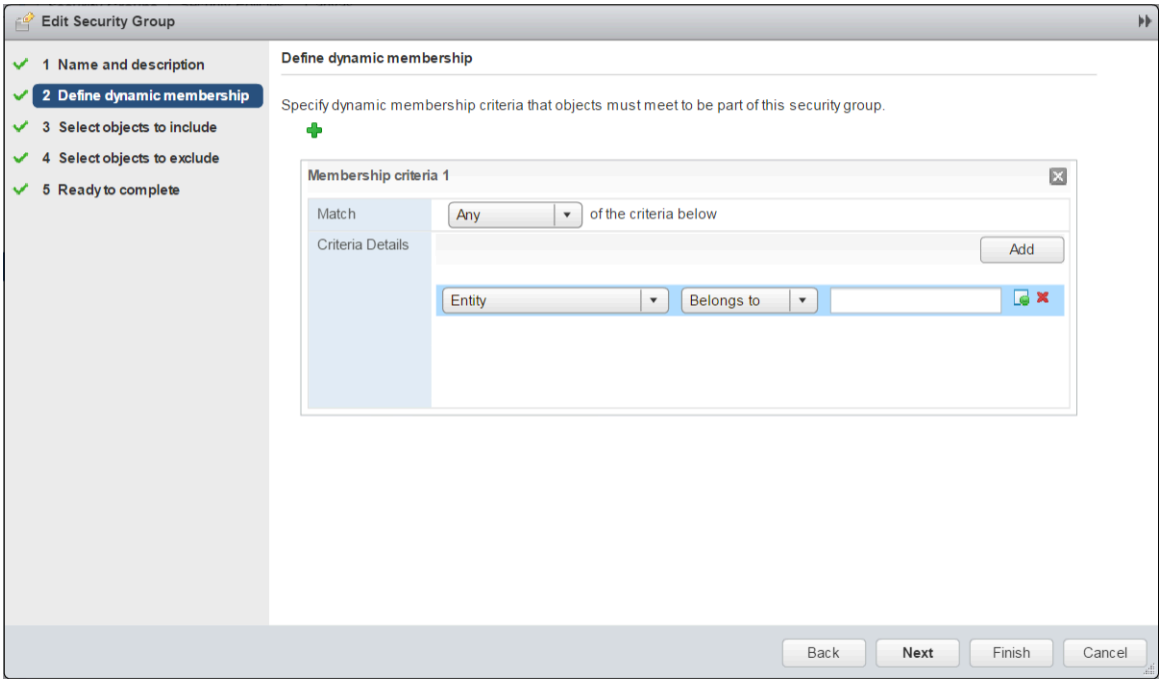

7. Click **Finish**.
8. Repeat this procedure for each security group that should have its traffic redirected to the VM-Series firewall.

Enable Large Receive Offload

Large receive offload (LRO) is a technique for increasing the inbound throughput on high-bandwidth network connections by decreasing CPU overhead. Without LRO, the firewall drops packets larger than the configured maximum transmission unit MTU, which is a maximum of 9216 bytes when the firewall is enabled for jumbo frames. With LRO enabled, the firewall accepts packets up to 64KB in size and the does not drop packets larger than the configured MTU. Instead, it segments the larger packets into smaller chunks of 9000 bytes. For example, if the VM1 sends a 64KB packet to VM2 and the packet is divided into eight segments.

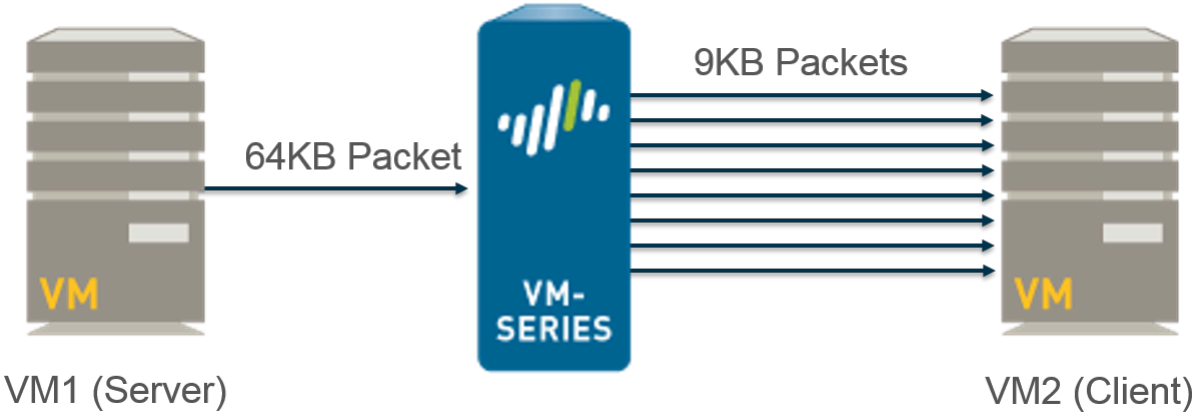

 ©2024 Palo Alto Networks, Inc.

LRO is disabled by default on new NSX-V deployments and upon upgrade. You can enable or disable LRO and view the LRO status on through the CLI. Enabling LRO on the VM-Series firewall automatically enables jumbo frames. Additionally, LRO and TCP Segmentation Offload (TSO) must be enabled on VMXNET3 network adapter on the VM-Series firewall host machine.

STEP 1 | Verify that large receive offload and TCP segmentation offload is enabled on the host.

For information about LRO and TSO on the host machine, see theVMware vSphere documentation.

1. Log in to vSphere and navigate to your host machine.
2. Select **Manage** > **Settings** > **System** > **Advanced System Settings.**
3. Locate the following parameters and verify that their value is set 1. A 1 indicates that the parameter is enabled on the VMXNET3 adapter.
   - For LRO—Net.Vmxnet3HwLRO
   - For TSO—Net.UseHwTSO and Net.UseHwTSO6

STEP 2 | Enable LRO on the VM-Series firewall.

1. Access the firewall CLI.
2. Use the following command to enable LRO:

```
admin@PA-VM> set system setting lro enable
```

3. Reboot the firewall using the following command:

```
> request restart system
```

4. Verify the LRO is enabled with the following command:

```
admin@PA-VM> show system setting lro
Device LRO mode:                        on
Current device mtu size:                9192
```

You can disable LRO using the command ***set system setting lro disable****.*

Create Security Groups and Steering Rules

The following topics describe how to create security groups and policies to steer traffic to the VM-Series firewall. Follow the link below that matches your deployment process—Security Centric or Operations Centric.

- Create Security Groups and Steering Rules in a Security Centric Deployment
- Create Security Groups and Steering Rules in an Operations Centric Deployment

 ©2024 Palo Alto Networks, Inc.

Create Security Groups and Steering Rules in a Security Centric Deployment

The following topics describe how to create policies on Panorama to steer traffic to the VM-Series firewall. In order for the VM-Series firewall to secure traffic, you must complete the following tasks:

- Set Up Dynamic Address Groups on Panorama
- Create Steering Rules on Panorama

Set Up Dynamic Address Groups on Panorama

A security group is a logical container that assembles guests across multiple ESXi hosts in the cluster. When you create a dynamic address group that meets the right criteria and commit your changes, a corresponding security group is created on the NSX-V Manager. Creating security groups are required to manage and secure the guests; to understand how security groups enable policy enforcement, see Policy Enforcement using Dynamic Address Groups.

STEP 1 | Configure a dynamic address group for each security group required for your deployment.

Shared dynamic address groups are not supported on the VM-Series for VMware NSX-V.

1. Select **Objects** > **Address Groups**.
2. Verify that you are configuring the dynamic address groups in a device group associated with an NSX-V service definition.
3. Click **Add** and enter a **Name** and **Description** for the address group.
4. Select **Type** as **Dynamic**.
5. Define the match criteria.

   For the dynamic address group to become a security group in NSX-V Manager, the match criteria string must be enclosed in single quotes with the prefix _nsx_ followed by the exact name of the Address Group. For example, ***'_nsx_PAN_APP_NSX'***.
6. Repeat this process for each security group you require.

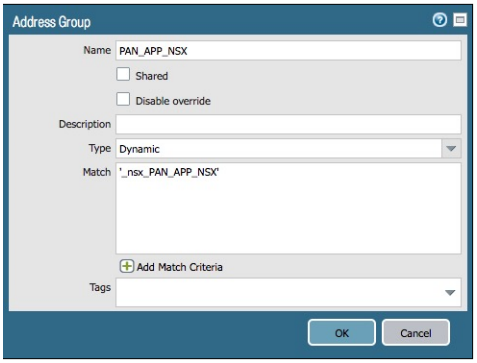

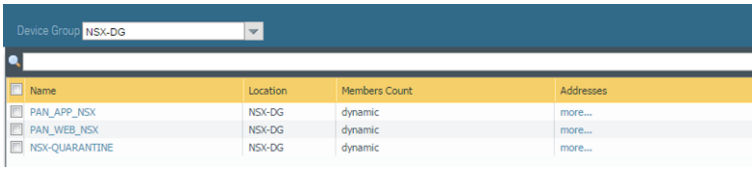

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Verify that the corresponding security groups are created on the NSX-V Manager.

1. Select **Network and Security** > **Service Composer** > **Security Groups**.
2. Verify that your dynamic address groups appear as security groups on the Security Groups list. Each security group is prefixed with your service definition followed by an underscore and the dynamic address group name.

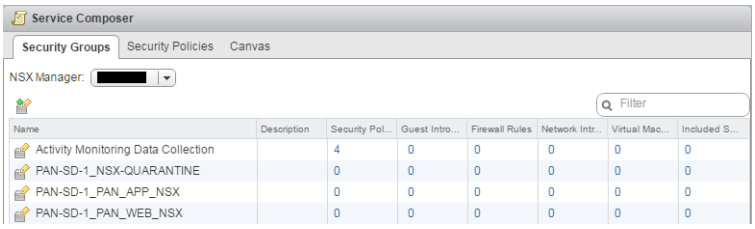

Create Steering Rules on Panorama

Do **not** apply the traffic redirection policies unless you understand how rules work on the NSX-V Manager as well as on the VM-Series firewall and Panorama. The default policy on the VM-Series firewall is set to *deny all* traffic, which means that all traffic redirected to the VM-Series firewall will be dropped. To create policies on Panorama and push them to the VM-Series firewall, see Apply Policies to the VM-Series Firewall.

Create security policy rules in the associated device group. For each security rule set the Rule Type to Intrazone, select one zone in the associated template stack, and select the dynamic address groups as the source and destination. Creating a qualifying security policy in Panorama helps in the creation of a corresponding steering rule on NSX-V Manager upon steering rule generation and commit in Panorama.

STEP 1 | Create security policy.

1. In Panorama, select **Policies** > **Security** > **Pre Rules**.
2. Verify that you are configuring the dynamic address groups in a device group associated with an NSX-V service definition.
3. Click **Add** and enter a **Name** and **Description** for your security policy rule.
4. Set the Rule Type to **intrazone (Devices with PAN-OS 6.1 or later)**.
5. In the Source tab, set the source zone to the zone from the template stack associated with the service definition. Then select a dynamic address group (NSX-V security group) you created previously as the Source Address. Do not add any static address groups, IP ranges, or netmasks as a Source Address.
6. In the Destination tab, Panorama does not allow you to set a destination zone because you set the rule type to intrazone. Then select a dynamic address group (NSX-V security group) you created previously as the Destination Address. Do not add any static address groups, IP ranges, or netmasks as a Destination Address.
7. Click **OK**.
8. Repeat steps 1 through 7 for each steering rule you require.
9. **Commit** your changes.

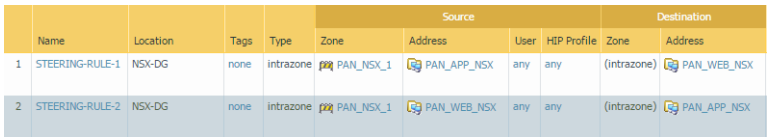

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Generate steering rules.

Panorama generates a steering rule for each qualifying security policy rule.

1. Select **Panorama** > **VMware** > **NSX-V** > **Steering Rules**.
2. Select **Auto-Generate Steering Rules**.

   Panorama will populate the list of steering rules based on qualified security policy rules in the device group attached in the service definition.

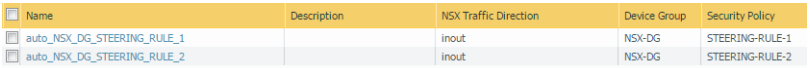

3. (Optional) Modify the NSX Traffic Direction and add NSX-V Services to a Steering Rule.

   By default, the NSX Traffic Direction is set to **inout** and no NSX-V Services are selected. When no NSX-V Services are specified, any type of traffic is redirected to the VM-Series firewall.

   1. Select the auto-generated steering to be modified.
   2. To change the traffic direction, select the direction from the **NSX Traffic Direction** drop-down.
   3. Click **Add** under NSX Services and choose a service from the **Services** drop-down. Repeat this step to add additional services.
   4. Click **OK**.
4. If you deleted any steering rules, click **Auto-Generate Steering Rules** before committing your changes.
5. **Commit** your changes.

STEP 3 | Verify that the corresponding traffic steering rules were created on the NSX-V Manager.

1. Select **Network and Security** > **Firewall** > **Configuration** > **Partner Security Services**.
2. Confirm that the traffic steering rules your created on Panorama are listed.

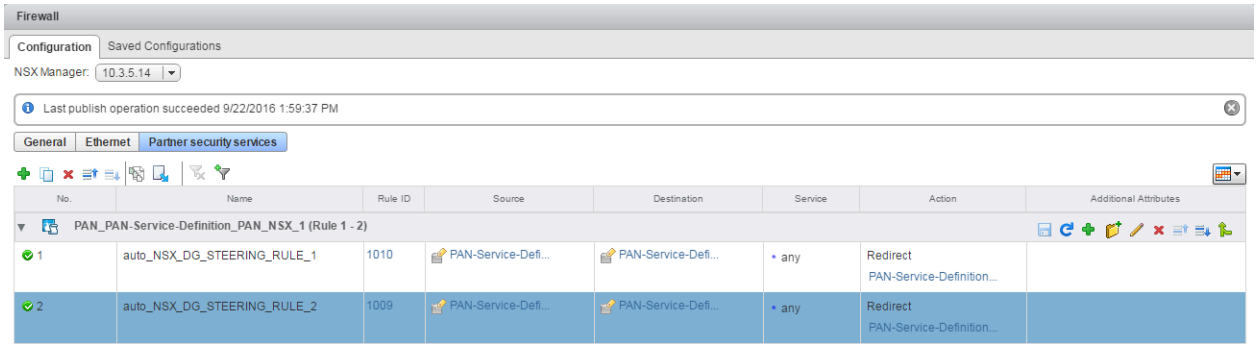

Create Security Groups and Steering Rules in an Operations Centric Deployment

In an operations-centric deployment, you create security groups and traffic redirection rules on the NSX-V Manager instead of Panorama. Then your security rules configured on Panorama enforce the traffic redirected to the VM-Series firewall. Complete the following tasks when deploying the VM-Series firewall for NSX-V in an operations-centric deployment:

- Set Up Security Groups on the NSX-V Manager

 ©2024 Palo Alto Networks, Inc.

- Create Steering Rules on NSX-V Manager

Set Up Security Groups on the NSX-V Manager

A security group is a logical container that assembles guests across multiple ESXi hosts in the cluster. Creating security groups makes it easier to manage and secure the guests. The VM-Series firewall can secure static VM membership, using IP Sets that include subnets and ranges, and dynamic VM membership, using tags. When using an IP Set as the membership criteriaTo understand how security groups enable policy enforcement, see Policy Enforcement using Dynamic Address Groups.

STEP 1 | Log in to the vSphere user interface.

STEP 2 | Select **Networking and Security** > **Service Composer** > **Security Groups**, and add a **New Security Group**.

STEP 3 | Add a **Name** and **Description**. This name will display in the match criteria list when defining dynamic address groups on Panorama.

STEP 4 | Select the guests that constitute the security group. You can either add members dynamically or statically. You can **Define Dynamic Membership** by matching on security tags (recommended), or statically by adding IP Sets under **Select the Objects to Include**. In the following screenshot, the guests that belong to the security group are selected using the **Objects Type: Virtual Machine** option.

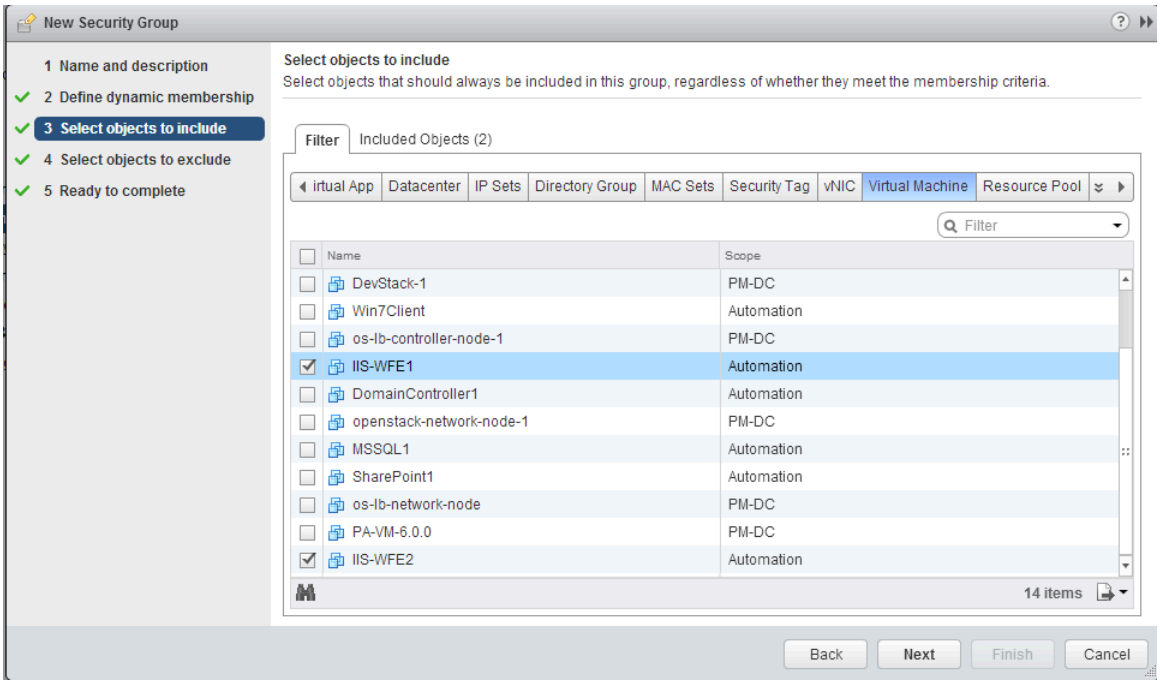

STEP 5 | Review the details and click **OK** to create the security group.

Create Steering Rules on NSX-V Manager

Do not apply the traffic redirection policies unless you understand how rules work on the NSX-V Manager as well as on the VM-Series firewall and Panorama. The default policy on the VM-Series firewall is set to deny all traffic, which means that all traffic redirected to the VM-Series firewall

 ©2024 Palo Alto Networks, Inc.

will be dropped. To create policies on Panorama and push them to the VM-Series firewall, see Apply Security Policies to the VM-Series Firewall.

STEP 1 | Select **Networking and Security** > **Service Composer** > **Security Policies** and click **Create Security Policy** ( ).

STEP 2 | Add a rule **Name**.

STEP 3 | Add a network introspective service.

1. Select **Network Introspection Service** and click the green plus icon.
2. **Name** the network introspection service and add a **Description**.
3. Select **Redirect to Service** under Action.
4. Select your service definition under Service Name.
5. Select you service profile under Profile.
6. Select a **Source** and a **Destination**. By default, traffic source is set to Policy's Security Groups. This option dynamically includes all security groups where this policy is applied. Alternatively, you can choose to have traffic from any source redirected to the firewall or specify certain security groups. However, vSphere requires that Source or Destination (or bother) be set Policy's Security Group. If you select Any or specific security groups for Destination, then Source must be set to Policy's Security Group.
7. (Optional) Select specific network services to be redirected to the firewall. If you choose any service or services, all other traffic will not be redirect to the firewall.
8. Click **OK**.
9. Repeat steps 1 through 6 to add additional network introspection services.
10. Click **Finish** to save your configuration.

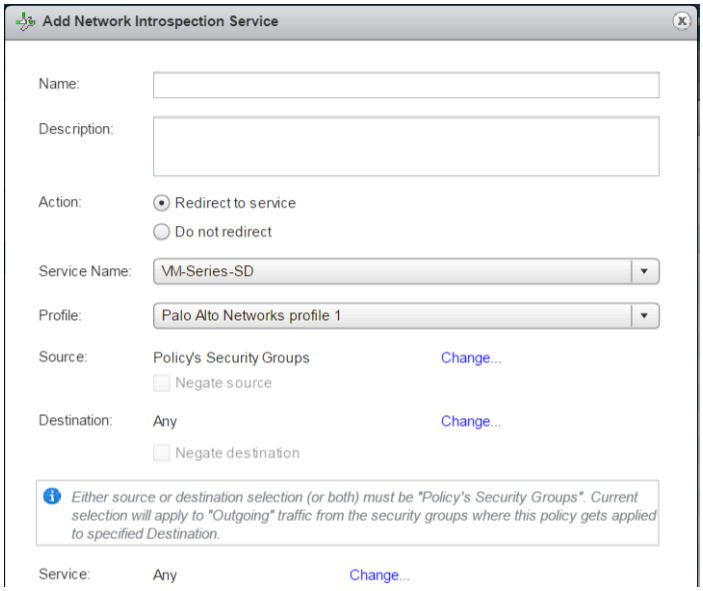

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Apply redirection policy to security groups.

1. Highlight a security policy by clicking it.
2. Select **Networking and Security** > **Service Composer** > **Security Policies** and click Apply Security Policy ( ).
3. Apply the redirection rules by checking all appropriate zones.
4. Click **OK**.

Apply Security Policies to the VM-Series Firewall

Now that you have created the steering rules on Panorama and pushed them to the NSX-V Manager, you can now use Panorama for centrally administering policies on the VM-Series firewalls.

To manage centralized policy, attach the dynamic address group as a source or destination address in security policy and push it to the firewalls; the firewalls can dynamically retrieve the IP addresses of the virtual machines that are included in each security group to enforce compliance for traffic that originates from or is destined to the virtual machines in the specified group.

STEP 1 | Log in to Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | (Operations-centric deployments only) Create dynamic-address groups.

Skip this step for security-centric deployments. If you are performing a security-centric deployment, you have already created dynamic-address groups.

After creating the security redirection rules on the NSX-V Manager, the names of the security groups that are referenced in security policy will be available on Panorama.

Shared dynamic address groups are not supported on the VM-Series for VMware NSX-V.

1. Select **Objects** > **Address Groups**.
2. Select the Device Group you created for managing your VM-Series on NSX-V firewall from the **Device Group** drop-down.
3. Click **Add** and enter a **Name** and **Description** for the dynamic address group.
4. Select **Type** as **Dynamic**.
5. Add Match Criteria to your dynamic address group.

   Some browser extensions may block API calls between Panorama and NSX-V which prevents Panorama from receiving match criteria. If Panorama displays no match criteria and you are using browser extensions, disable the extensions and Synchronize Dynamic Objects to populate the tags available to Panorama.
6. Click **Add Match Criteria**.
7. Select the **And** or **Or** operator and click the plus (+) icon next to the security group name to add it to the dynamic address group.

   The security groups that display in the match criteria dialog are derived from the groups you defined on the Distributed Firewall Partner Security Services or on the Service Composer on the NSX-V Manager. Only the security groups that are referenced in the security policies and from which traffic is redirected to the VM-Series firewall are available here.

 ©2024 Palo Alto Networks, Inc.

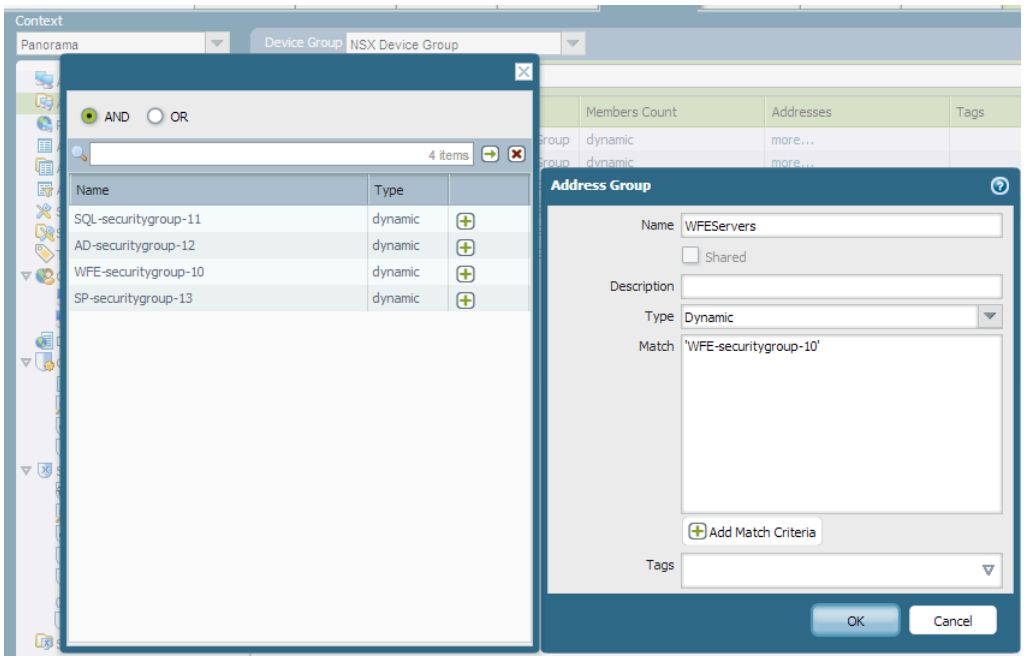

8. Click **OK**.
9. Repeat these steps to create the appropriate number of dynamic address groups required for your deployment.
10. **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create security policy rules.

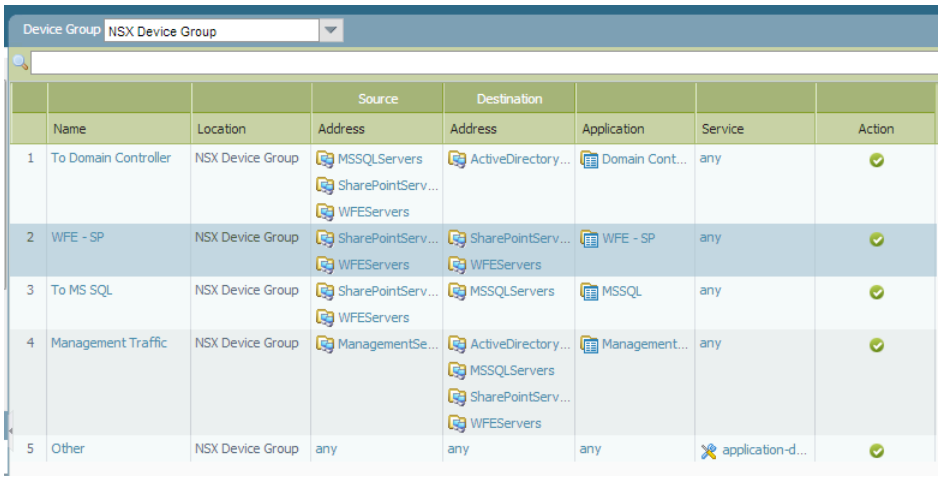

1. Select **Policies** > **Security** > **Prerules**.
2. Select the **Device Group** 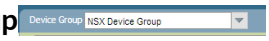

that you created for managing the VM-Series firewalls for NSX-V in Register the VM-Series Firewall as a Service on the NSX-V Manager.
3. Click **Add** and enter a **Name** and a **Description** for the rule. In this example, the security rule allows all traffic between the WebFrontEnd servers and the Application servers.
4. Select the **Source Zone** and **Destination Zone**. The zone name must be the same in both columns.
5. For the **Source Address** and **Destination Address**, select or type in an address, address group or region. In this example, we select an address group, the Dynamic address group you created previously.

6. Select the **Application** to allow. In this example, we create an **Application Group** that includes a static group of specific applications that are grouped together.
   1. Click **Add** and select **New Application Group**.
   2. Click **Add** to select the application to include in the group. In this example, we select the following:
   3. Click **OK** to create the application group.

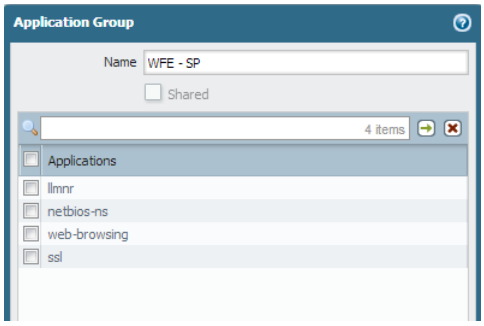

 ©2024 Palo Alto Networks, Inc.

7. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles for antivirus, anti-spyware, and vulnerability protection, under Profiles.
8. Repeats the steps above to create the pertinent policy rules.
9. Click **Commit**, select Commit Type as **Panorama**. Click **OK**.

STEP 4 | Apply the policies to the VM-Series firewalls for NSX-V.

1. Click **Commit**, and select Commit Type **Device Groups**.
2. Select the device group, NSX-V Device Group in this example and click **OK**.
3. Verify that the commit is successful.

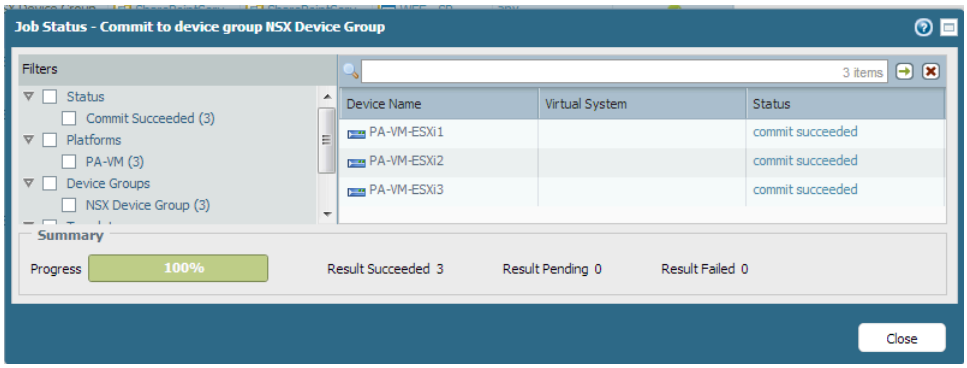

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Validate that the members of the dynamic address group are populated on the VM-Series firewall.

1. From Panorama, switch device context to launch the web interface of a firewall to which you pushed policies.

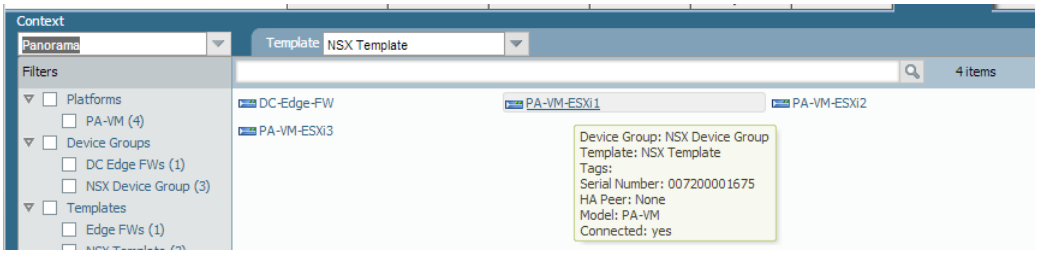

2. On the VM-Series firewall, select **Policies** > **Security**, and select a rule.
3. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.

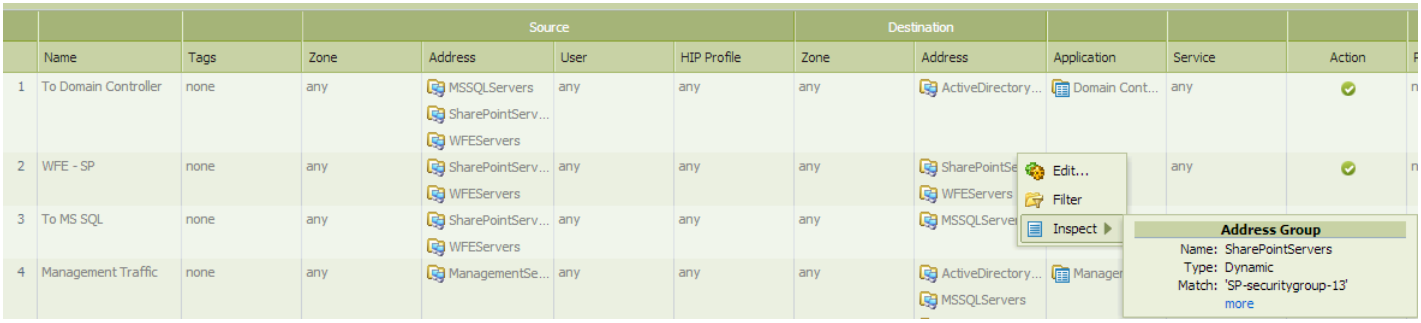

4. Click the **more** link and verify that the list of registered IP addresses is displayed.

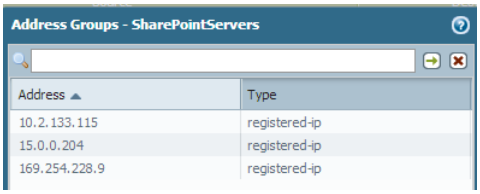

Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

STEP 6 | (Optional) Use template to push a base configuration for network and device configuration such as DNS server, NTP server, Syslog server, and login banner.

Refer to the Panorama Administrator's Guide for information on using templates.

STEP 7 | Create a Zone Protection profile and attach it to a zone.

A zone protection profile provides flood protection and has the ability to protect against port scanning, port sweeps and packet-based attacks. It allows you to secure intra-tier and inter-tier

 ©2024 Palo Alto Networks, Inc.

traffic between virtual machines within your data center and traffic from the Internet that is destined to the virtual machines (workloads) in your data center.

1. Select your **Template**.
2. Select **Network** > **Network Profiles** > **Zone Protection** to add and configure a new profile.
3. Select **Network** > **Zones**, click the default-zone listed and select the profile in the **Zone Protection Profile** drop down.

STEP 8 | Create a DoS Protection profile and attach it to DoS Protection policy rule.

1. Select your **Device Group**.
2. Select **Objects** > **Security Profiles** > **DoS Protection** to add and configure a new profile.
   - A classified profile allows the creation of a threshold that applies to a single source IP. For example, you can configure a max session rate for an IP address that matched the policy, and then block that single IP address once the threshold is triggered.
   - An aggregate profile allows the creation of a max session rate for all packets matching the policy. The threshold applies to new session rate for all IP addresses combined. Once the threshold is triggered it affects all traffic that matches the policy.
3. Create a new DoS Protection policy rule in **Policy** > **DoS Protection,** and attach the new profile to it.

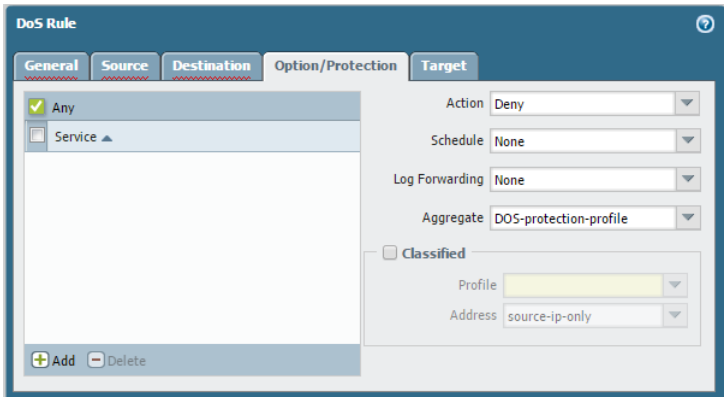

Steer Traffic from Guests that are not Running VMware Tools

VMware Tools contains a utility that allows the NSX-V Manager to collect the IP address(es) of each guest running in the cluster. NSX-V Manager uses the IP address as a match criterion to steer traffic to the VM-Series firewall. If you do not have VMware tools installed on each guest, the IP address(es) of the guest is unavailable to the NSX-V Manager and traffic cannot be steered to the VM-Series firewall.

The following steps allow you to manually provision guests without VMware Tools so that traffic from each of these guests can be managed by the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create an IP set that includes the guests that need to be secured by the VM-Series firewall. This IP set will be used as the source or destination object in an NSX-V distributed firewall rule in Step 2 below.

1. Select **NSX Managers** > **Manage** > **Grouping Objects** > **IP Sets**.
2. Click **Add** and enter the IP address of each guest that does not have VMware tools installed, and needs to be secured by the VM-Series firewall. Use commas to separate individual IP addresses; IP ranges or subnets are not valid.

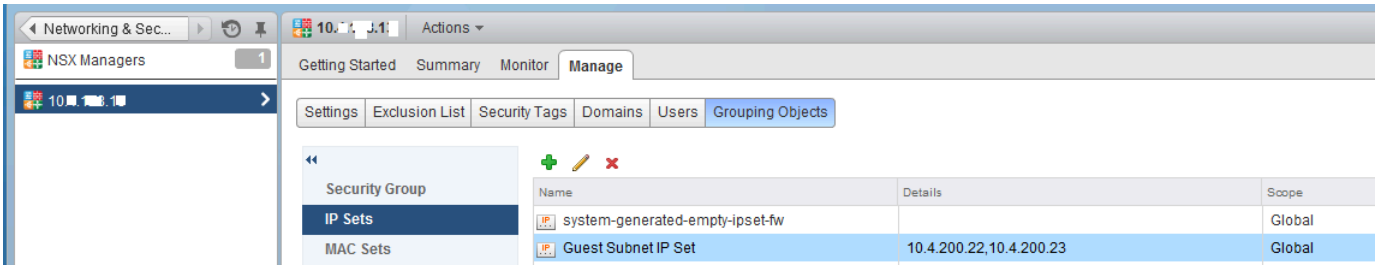

STEP 2 | Attach the IP sets to the Security Groups on NSX-V, to enforce policy.

1. Select **Networking and Security** > **Service Composer** > **Security Groups**.
2. Select **Select objects to include** > **IP Sets**, add the IP set object to include.

What is Multi-NSX Manager Support on the VM-Series for NSX-V?

Multi-NSX Manager support on the VM-Series firewall for NSX-V allows you to connect a single Panorama to multiple NSX-V Managers running individual vCenter servers. Using a single Panorama allows you to manage common objects and policies and synchronize them across multiple vCenter servers. You can now configure and manage multiple NSX-V Managers in a single location, eliminating the need to replicate common configuration many times on multiple Panorama servers.

- Plan Your Multi-NSX Deployment
- Deploy the VM-Series Firewall in a Multi-NSX Manager Environment

Plan Your Multi-NSX Deployment

You must carefully plan your device group hierarchy and template stacks and consider how they interact with the other components needed for deployment. Service definitions reference device groups and template stacks and push that information to the firewalls in the related ESXi clusters.

- **Configure your device groups**—Devices groups are logical units that group firewalls based on common aspects that require similar policy configurations. Each service definition requires a device group and each device group can only be referenced in one service definition.

  A device group inherits policy rules and object settings from device groups above it in the device group hierarchy. This allows you to configure common or shared settings in parent device groups and unique settings in child or grandchild device groups. By default, Panorama has a Shared device group and any configuration in the shared device group is pushed to all

 ©2024 Palo Alto Networks, Inc.

device groups. When configuring any policy rules or object settings, confirm that you have selected the right device group.

See Managing Device Groups in the Panorama Administrator's Guide for information on configuring and managing device groups.

- **Configure your template stacks**—A template stack contains settings that enable a firewall to connect to your network, such as interface and zone configurations. Each service definition requires a template stack and each template stack can only be referenced in one service definition.

  When assigning a template stack to a service definition, consider the priority of the templates in the stack to ensure that the right configuration is pushed to the correct firewalls. If the templates in a stack contain overlapping configuration, the template with higher priority takes precedence and the same setting in lower templates are ignored. Therefore, ensure that template configuration unique to an NSX-V Manager is given higher priority in the template stack assigned to that NSX-V Manager's service definition.

  See Manage Templates and Template Stacks in the Panorama Administrator's Guide for information on configuring and managing template stacks.

- **Create your service definition**—A service definition specifies the configuration for the VM-Series firewalls on each host in the ESXi cluster. Each individual NSX-V manager configuration requires at least one service definition. A service manager can have multiple service definitions but each service definition can only have one device group and one template stack. After a device group or template stack has been assigned to a service definition, you can no longer select that device group or template stack for future service definitions.

For example, in a disaster recovery deployment scenario, you would need to create identical device groups for each data center. Because all the policy rules and objects are the same for data centers, you can perform all you configuration in a single device group. However, you cannot use the same device group in two service definitions. To ensure that each data center gets the same policy rules, create a child device group for each data center under the device group with the common configuration. These child device groups do not need any configuration of their own because they inherit everything the VM-Series firewalls need from the parent device group. And because each data center is identical, configure your network settings in a template (Template 1). Create a template stack for each data center and assign Template 1 to each stack.

 ©2024 Palo Alto Networks, Inc.

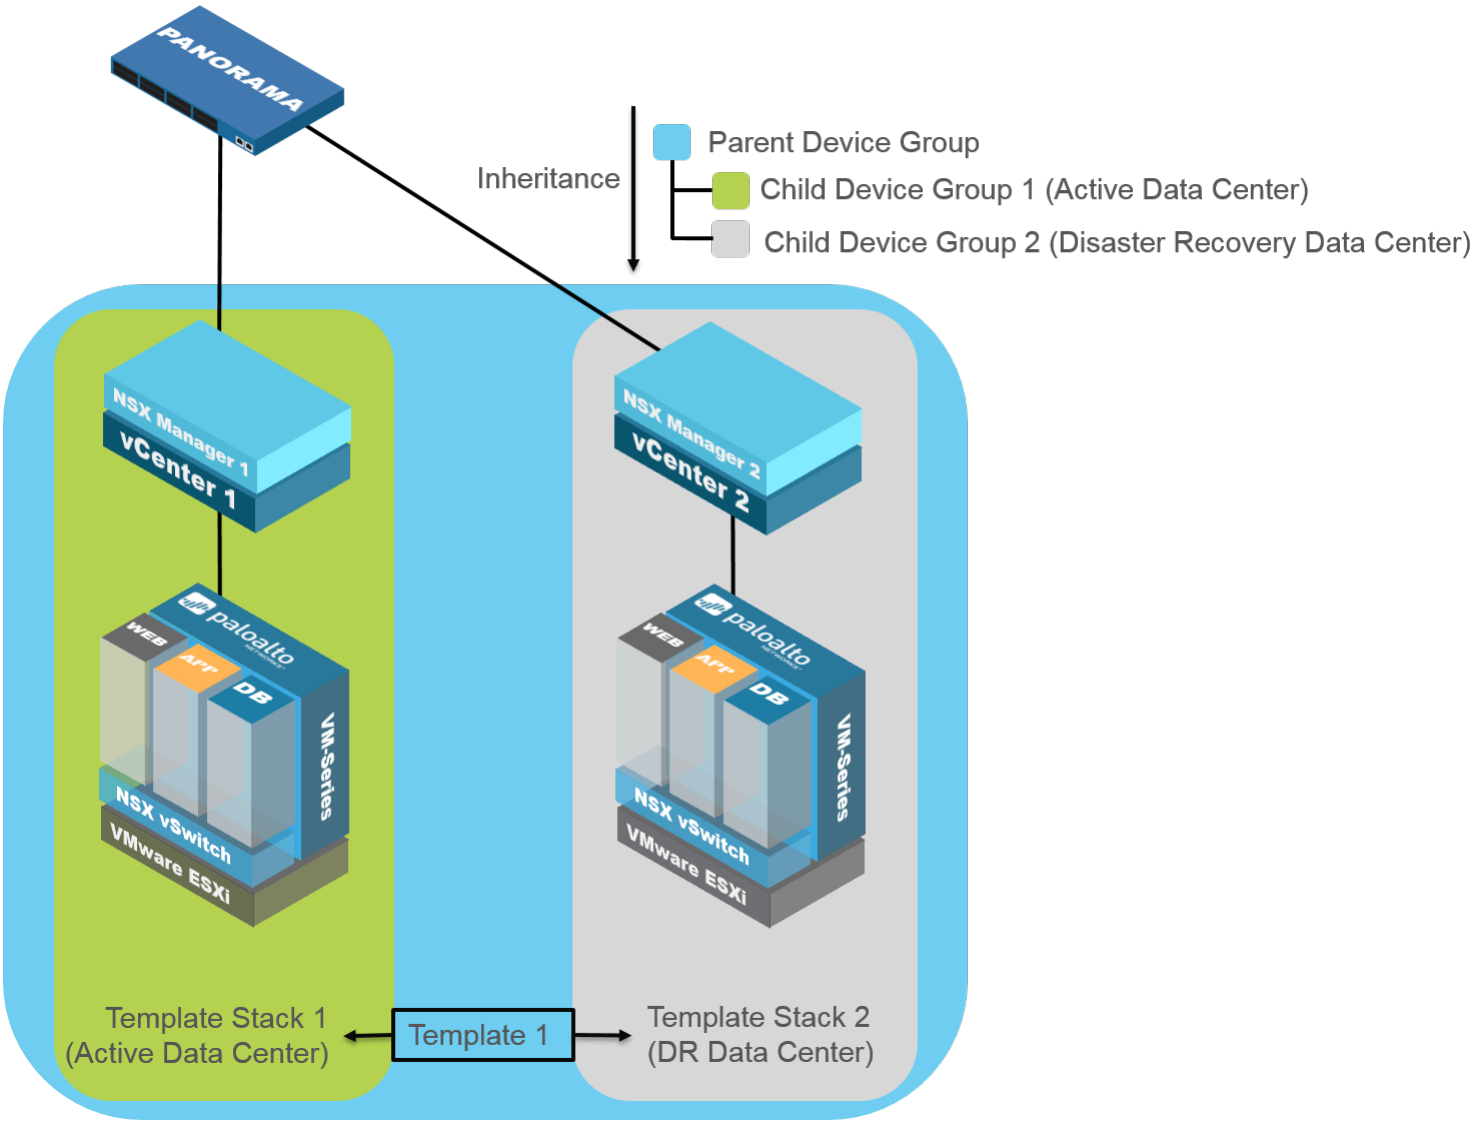

Deploy the VM-Series Firewall in a Multi-NSX Manager Environment

Whether you are deploying a single NSX-V Manager or a multi-NSX Manager environment, set up the connection between an NSX-V Manager and Panorama before you continue on to set up the next NSX-V Manager with Panorama.

STEP 1 | Install the VMware NSX Plugin version 2.0 as it allows you to connect up to 16 NSX-V Managers. This version of the plugin allows you to add more than one Service Manager to your VM-Series firewall for NSX-V configuration on Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enable Communication Between the NSX-V Manager and Panorama.

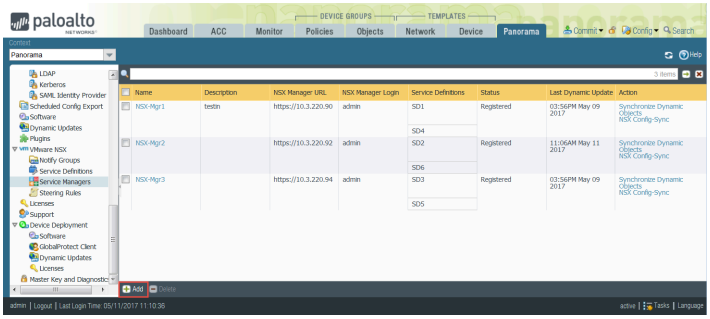

STEP 3 | Create Template(s), Template Stack(s), and Device Group(s) on Panorama. Device groups and template stacks push the security policy and network settings to the VM-Series firewalls in each ESXi cluster.

When configuring policy rules and objects, verify that you have selected the correct device group.

When configuring network and device settings, verify that you have selected the correct template stack.

STEP 4 | Create the Service Definitions on Panorama and attach them to the service manager. Each service definition can reference one device group and one template stack. Panorama supports up to 32 service definitions across all service managers.

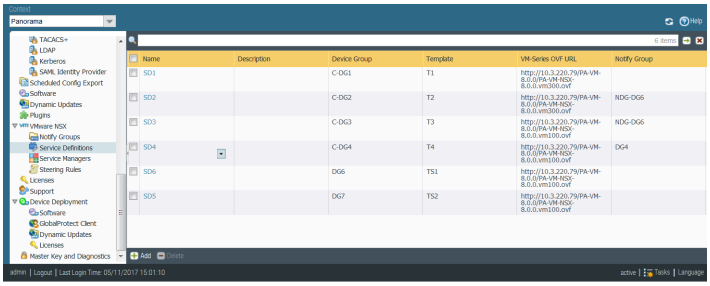

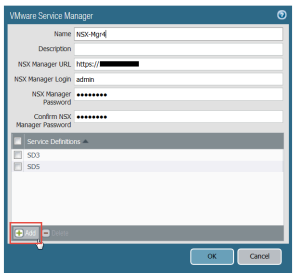

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Configure dynamic address groups or security groups and redirect traffic to the VM-Series firewall.

- For security-centric deployments Set Up Dynamic Address Groups on Panorama and Create Steering Rules on Panorama.
- For operations-centric deployments Set Up Security Groups on the NSX-V Manager and Create Steering Rules on NSX-V Manager.

Verify that you have selected the correct device group so the right steering rules are sent to the corresponding NSX-V Manager.

STEP 6 | Deploy the Palo Alto Networks NGFW Service on each ESXi cluster by using the relevant service definitions.

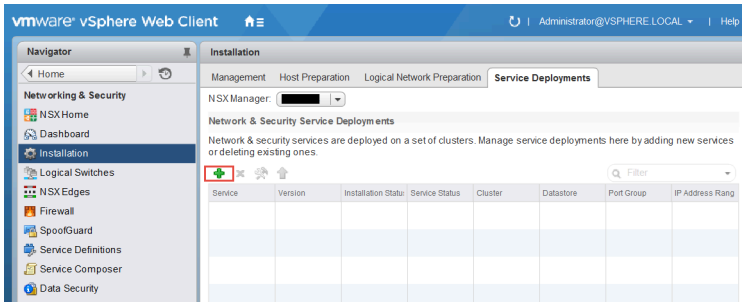

STEP 7 | Repeat this process for each NSX-V Manager.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Managers** and click **Add**.
2. Enable Communication Between the NSX-V Manager and Panorama.

Dynamically Quarantine Infected Guests

Threat and traffic logs in PAN-OS include the source or destination universally unique identifier (UUID) of guest VMs in your NSX-V deployment. This allows the VM-Series for NSX-V to support the tagging of guest VMs with NSX-V security tags. With the guest VMs' UUID now included in the log events, the firewall, based on the filtered log events, can tag the affected guest VM via NSX-V Manager API. This allows for automatic location of compromised VMs in the NSX-V environments. NSX-V can then put all associated UUIDs under policies to quarantine those VMs from the rest of the network.

Panorama includes predefined payload formats for threat and traffic logs in the HTTP Server Profile. These payload formats correspond to predefined security tags in NSX-V. When a guest VM is found in the threat or traffic logs, Panorama makes an API call to NSX-V Manager telling NSX-V Manager to tag the guest VM with the tag specified in the HTTP Server Profile. When the guest VM becomes tagged, NSX-V Manager dynamically moves the tagged guest VM into the quarantine security group, which places the guest VM into the quarantine dynamic address group.

STEP 1 | Confirm that you have content update version 636 or later installed on Panorama.

STEP 2 | Create a dynamic address group to be your quarantine dynamic address group.

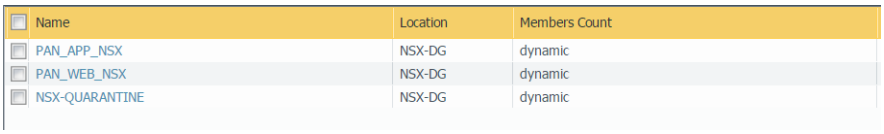

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create an HTTP Server Profile to send API calls to NSX-V Manager.

1. Select **Panorama** > **Server Profiles** > **HTTP** and **Add** a new HTTP Server Profile.
2. Enter a descriptive **Name**.
3. Select **Add** to provide the details of NSX-V Manager.
4. Enter a **Name** for NSX-V Manager.
5. Enter the **IP Address** of NSX-V Manager.
6. Select the Protocol (HTTP or HTTPS). The default Port is 80 or 443 respectively.
7. Select PUT under the HTTP Method column.
8. Enter the username and password for NSX-V Manager.

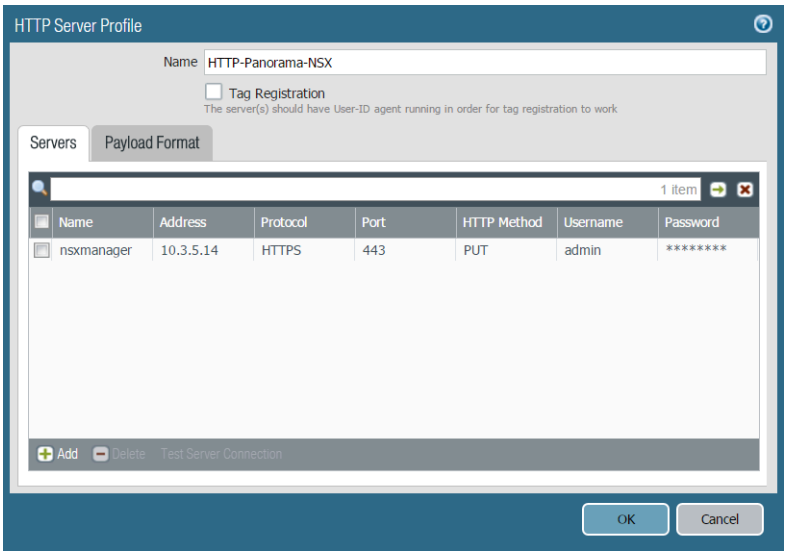

9. Select **Payload Format** and choose an NSX-V payload format from the Pre-defined Formats drop-down. This populates the URI Format, HTTP Headers, and Payload fields with the correct information to send the HTTP API call to NSX-V Manager. Additionally, the chosen format determines which security tag NSX-V Manager applies to infected guest VMs. In the example below, **NSX-V Anti-Virus Threat High** is selected which

 ©2024 Palo Alto Networks, Inc.

corresponds to the **ANTI_VIRUS.VirusFound.threat=high** security tag on NSX-V Manager.

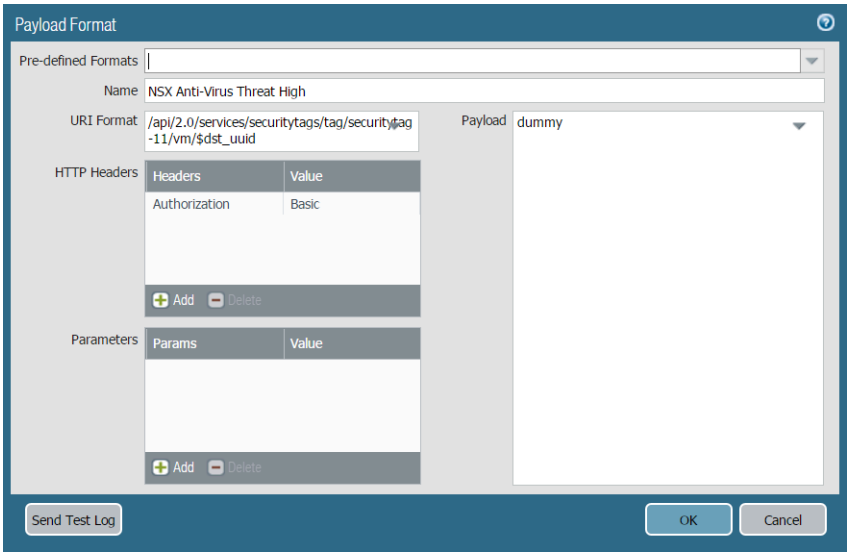

STEP 4 | Define the match criteria for when Panorama will forward logs to the NSX-V Manager, and attach the HTTP server profile to use.

1. Select **Panorama** > **Collector Groups** > **Collector Log Forwarding** for Threat or Traffic logs.
2. Click **Traffic** or **Threat** and **Add**.
3. Enter a descriptive name for the new log settings.
4. (Optional) Under Filter, you can add filters such as severity to narrow the logs that are forwarded to NSX-V Manager. If All Logs is selected, all threat or traffic logs that meet the criteria set in the HTTP Server profile are sent to NSX-V Manager.
5. Click **Add** under HTTP and select the HTTP Server Profile configured in Step 3.
6. Click **OK**.

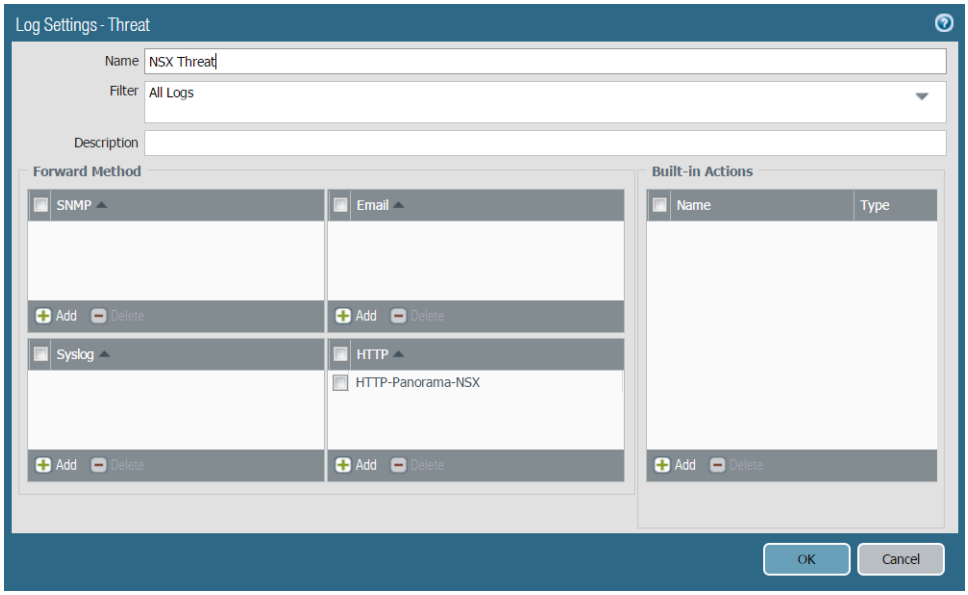

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Configure an NSX-V server certificate for Panorama to forward logs to NSX-V manager.

1. Select **Panorama** > **Certificate Management** > **Certificates**.
2. Create a root CA certificate with CN=IP address of Panorama.
3. Create a signed certificate with CN=IP address of NSX-V Manager.
4. Export the root CA certificate in PEM format without a private key.
5. Export the signed certificate in PEM format with a private key.

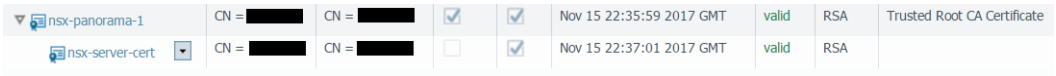

6. Using a tool such as OpenSSL, concatenate the exported certificates into a single PEM file for upload to NSX-V manager. Use the following commands in OpenSSL to complete this step.

```
cat cert_NSX_Root_CA.crt
cert_NSX_Signed1.pem > cert_NSX_cert_chain.pem
openssl pkcs12 -export -in cert_NSX_cert_chain.pem -out
 cert_NSX_cert.p12
```

7. Log in to NSX-V Manager and select **Manage Appliance Settings** > **SSL Certificates** > **Upload PKC#12 Keystore**. Click **Choose File**, locate the p12 file you created in the previous step, and click **Import**.

STEP 6 | Associate a security group with a security tag in vCenter.

1. Log in to vCenter.
2. Select **Networking & Security** > **Service Composer** > **Security Groups**.
3. Select a security group that is counterpart to the quarantine dynamic address group you created previously and click **Edit Security Group**.

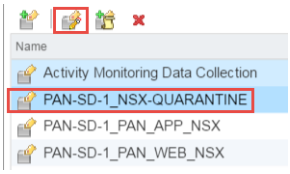

4. Select Define dynamic membership and click the + icon.
5. Click **Add**.
6. Set the criteria details to Security Tag Contains and then enter the NSX-V security tag that corresponds to the NSX payload format you chose in 3. Each of the predefined NSX-V payload formats corresponds to an NSX-V security tag. To view the NSX-V

 ©2024 Palo Alto Networks, Inc.

security tags in NSX-V, select **Networking & Security** > **NSX Managers** > **NSX Manager IP** > **Manage** > **Security Tags**.

In this example, **NSX Anti-Virus Threat High** is used in the HTTP Server Profile so **ANTI_VIRUS.VirusFound.threat=high** is the NSX-V Security Tag that is used here.

7. Click Finish.

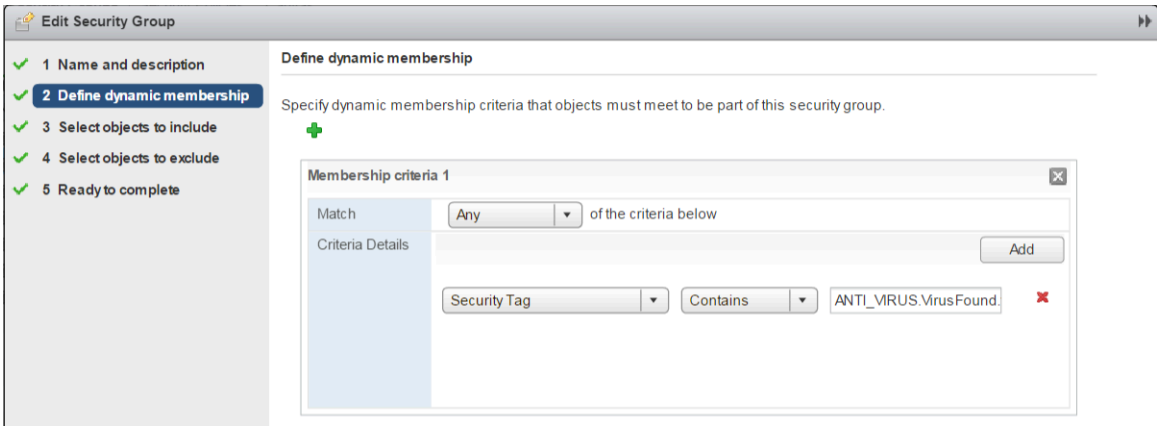

STEP 7 | After the guest VM is cleared for removal from quarantine, manually remove the NSX-V security tag from the guest VM in NSX-V.

1. Log in to vCenter.
2. Select **VMs and Templates** and choose the quarantined guest.
3. Select **Summary** > **Security Tags** > **Manage**.
4. Uncheck the security tag used by the quarantine security group and click OK.
5. Refresh the page and the quarantine security will no longer be listed under **Summary** > **Security Group Membership**.

Source and destination UUID fields in threat and traffic logs may be blank after a guest VM is removed from quarantine. This can occur when running NSX-V 6.2.3 or earlier or if NSX-V steering rules do not use the inout direction. You can resolve this by upgrading NSX-V to 6.2.4 or issue an NSX Config-sync under **Panorama** > **VMware** > **NSX-V** > **Service Manager** and reboot the PA-VM to resolve this issue.

Migrate Operations-Centric Configuration to Security-Centric Configuration

Complete the following procedure to migrate your Operations Centric configuration into Security Centric formats. This migration is not required. The VM-Series firewall for VMware NSX-V supports both styles of configuration. However, using both styles of configuration in the same deployment is not recommended.

STEP 1 | Upgrade Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Update the match criteria format in your dynamic address groups.

1. Select **Objects** > **Address Groups** and click the link name for your first dynamic address group.
2. Delete the existing match criteria entry.
3. Enter the new match criteria in the following format:

```
'_nsx_<dynamic-address-group-name>'
```

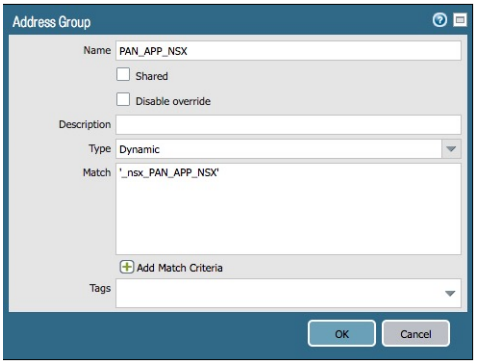

4. Click **OK**.
5. Repeat this process for each dynamic address group.

STEP 3 | Change security policy used as NSX-V steering rules to intrazone.

1. Select **Policies** > **Security** > **Pre Rules** and click the link name for your first security policy rule.
2. On the General tab, change the **Rule Type** to intrazone.

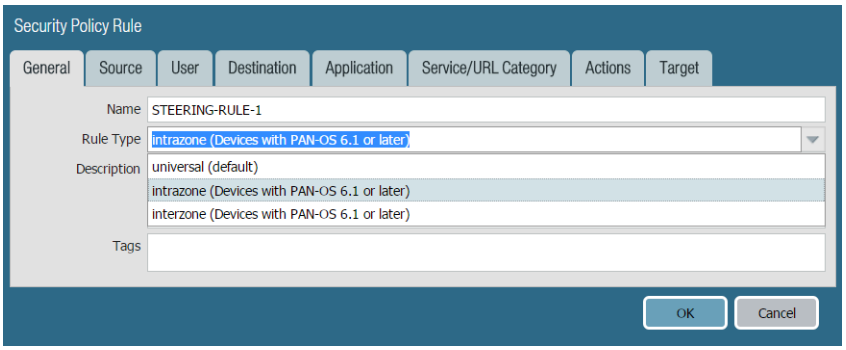

3. Click **OK**.
4. Repeat this process for each security policy rule.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Generate new steering rules.

1. Select **Panorama** > **VMware** > **NSX-V** > **Steering Rules**.
2. Click **Auto-Generate Steering Rules**.

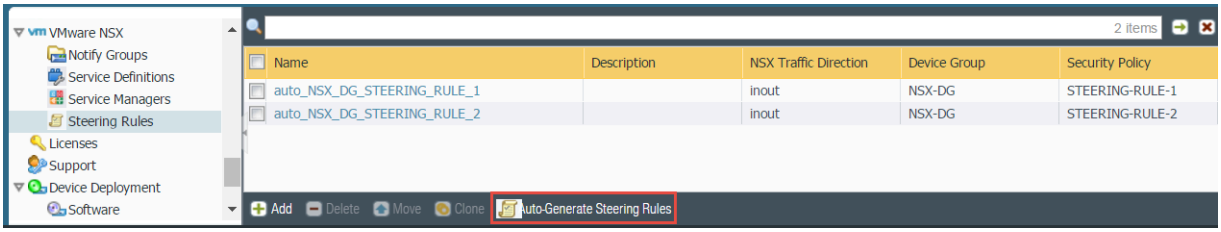

STEP 5 | **Commit** your changes.

When you commit your changes, Panorama pushes updates to NSX-V Manager.

1. Verify that NSX-V Manager created new security groups.
   1. Login to vCenter and select **Networking & Security** > **Security Groups**.
   2. The new security groups (mapped to the updated dynamic address groups) should appear in the following format:

      <service-definition-name> - <dynamic-address-group-name>

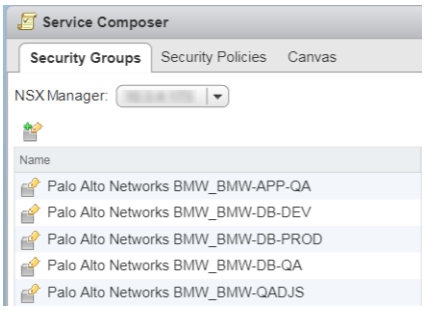

2. Verify that NSX-V Manager created new steering rules.
   1. Select **Networking & Security** > **Firewall** > **Configuration** > **Partner security services**.
   2. The new steering rules (mapped to the security policy rules you create on Panorama) are listed above the old steering rules.

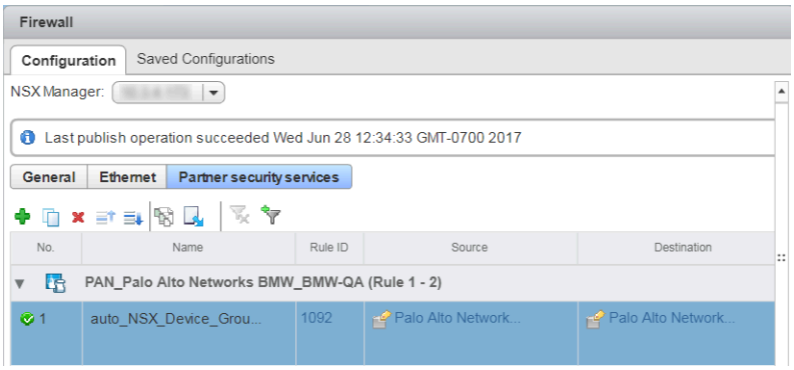

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Add match criteria to the newly created security groups to ensure that your VMs are placed in the correct security group.

There two ways to complete this task—recreate the match criteria from the old security group in the new security group or nest the old security group within the new security group.

To recreate the match criteria from the old security group, complete the following procedure.

1. Select **Network & Security** > **Service Composer** > **Security Groups**.
2. Click on a new security group and select **Edit Security Group**.
3. Select **Define dynamic membership** and click the plus icon.
4. Add the same match criteria in the corresponding old security group.
5. Repeat this process for each new security group.
6. Delete the old security groups.

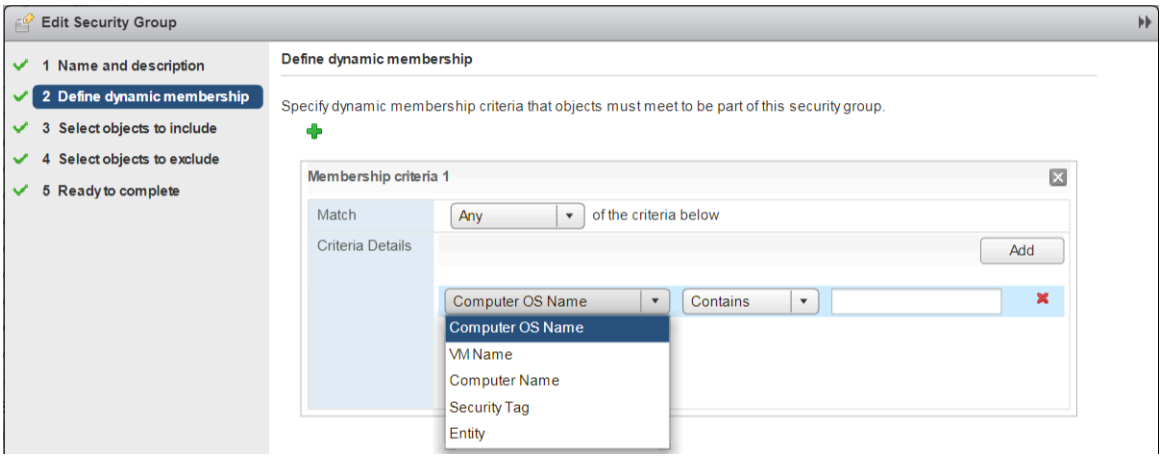

To nest the old security group within the new security group, complete the following procedure. In this method, VMs in the old security group are added to the new security group.

 ©2024 Palo Alto Networks, Inc.

Additionally, any new VM that meets the criteria of the old security group is automatically added to the new security group.

1. Select **Network & Security** > **Service Composer** > **Security Groups**.
2. Click on a new security group and select **Edit Security Group**.
3. Select **Select objects to include**.
4. Select the **Security Group** Object Type.
5. Choose the corresponding old security group under Available Objects and move it to Selected Objects by clicking the right arrow icon.
6. Click **Finish**.

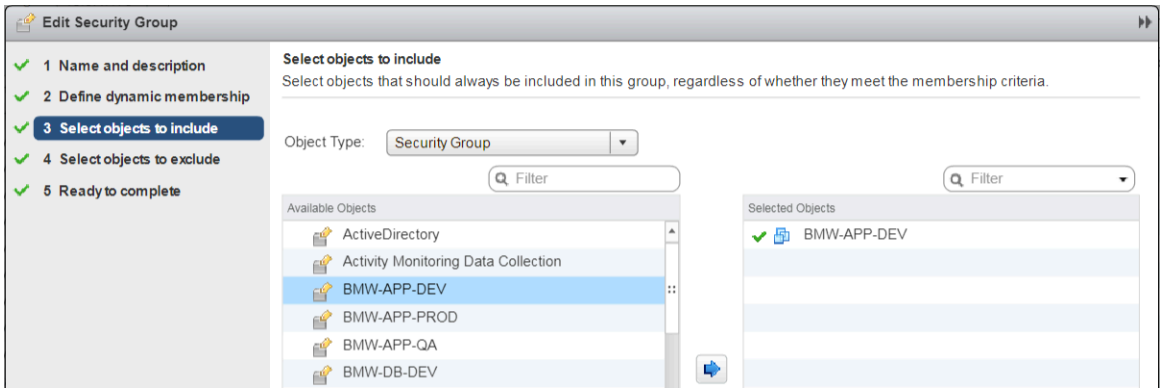

STEP 7 | Delete the old steering rules from vCenter.

1. Select **Networking & Security** > **Firewall** > **Configuration** > **Partner security services**.
2. Delete the old steering rules. Take care not to delete the Palo Alto Networks rules created by the Security-Centric workflow. These steering rule sections use the following naming convention.

   <service-definition-name> - <dynamic-address-group-name>

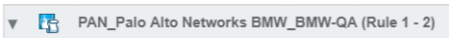

Add a New Host to Your NSX-V Deployment

When you add a new host to an existing cluster in your NSX-V environment, a new VM-Series firewall service is deployed on that new host. If you have DRS enabled and to ensure that the no traffic is lost while the VM-Series firewall is coming up, complete the following procedure. You will need to create a temporary cluster to house the new host and disable VMotion until the VM-Series firewall is ready.

STEP 1 | Create a temporary cluster and add your new host.

STEP 2 | Deploy the VM-Series Firewall service instance on your new host.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Disable vMotion on the virtual switches or distributed switches used by your new host.

- Virtual Switch
  1. Log in to vCenter Server and select your new host in the inventory.
  2. Select **Configuration** > **Networking** > **Virtual Switch**.
  3. Select the virtual switch that includes the VMkernal port group used by your host and click **Properties** > **Ports**.
  4. Select the port group configured for VMotion and click **Edit**.
  5. On the **General** tab, uncheck the **Enabled** check box for VMotion.
  6. Click **OK** and **Close**.
- Distributed Virtual Switch
  1. Log in to vCenter Server and select your new host in the inventory.
  2. Select **Configuration** > **Distributed Virtual Switch**.
  3. Select **Manage Virtual Adapters**.
  4. Choose the virtual adapter used by your host and click **Edit**.
  5. Uncheck **Use this virtual adapter for VMotion**.
  6. Click **OK**.

STEP 4 | Place your new host in maintenance mode.

STEP 5 | Move your new host to the existing cluster and exit maintenance mode. Wait for the VM-Series firewall to boot up completely. Because VMotion is disabled on the host, no other virtual machines will be moved to the host by DRS.

STEP 6 | When the VM-Series firewall is booted up completely, renable VMotion on the new host. Guest VMs can now move to the new host and security is intact with no traffic lost.

Use Case: Shared Compute Infrastructure and Shared Security Policies

This use case allows you to logically isolate traffic from two tenants that share an ESXi cluster and have a common set of security policies. In order to isolate traffic from each tenant you need to create a service definition with a template stack that includes two zones. Zone-based traffic separation makes it possible to distinguish traffic between virtual machines that belong to separate tenants, when it traverses through the firewall. The firewall is able to distinguish traffic between tenant virtual machines based on a service profiles and security groups created on the NSX-V Manager, which are available as match criteria in Dynamic Address Groups on the firewall. Therefore, even with overlapping IP addresses, you can segregate traffic from each tenant and secure each tenant's virtual machines using zone-base policy rules (source and destination zones must be the same) and dynamic address groups.

 ©2024 Palo Alto Networks, Inc.

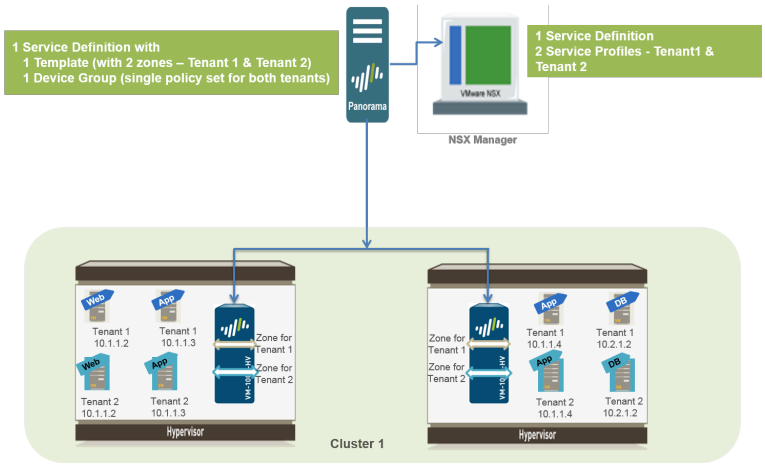

STEP 1 | Enable Communication Between the NSX-V Manager and Panorama.

This is one-time task and is required if you have not enabled access between the NSX-V Manager and Panorama.

STEP 2 | Create Template(s) and Device Group(s) on Panorama.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Templates** to add a template stack. This use case has a template stack named NSX-Template.
3. Select **Panorama** > **Device Groups** and add device group. This use case has a device group named NSX-DG.
4. Create two zones within the template stack. To isolate traffic for each tenant, you need two zones in this use case.
   1. Select **Network** > **Zones**.
   2. Select the correct template stack in the **Template** drop-down.
   3. Select **Add** and enter a zone **Name**. For example, **Tenant1**.
   4. Sets the interface **Type** to **Virtual Wire**.
   5. Click **OK**.
   6. Repeat the steps to add another zone, for example, **Tenant2**.
   7. Verify that the zones are attached to the correct template stack.

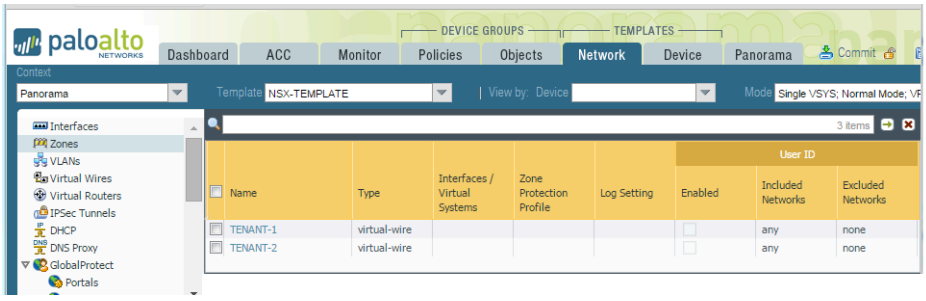

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create the Service Definitions on Panorama.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Definitions**.
2. Select **Add** and fill in the details.

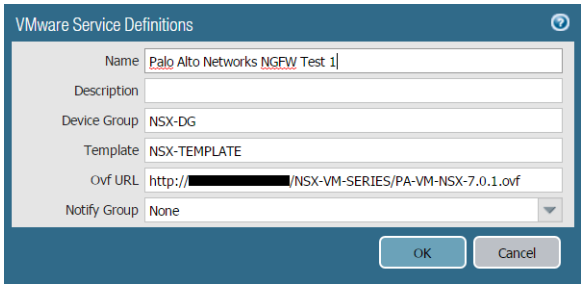

3. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 4 | Create Security Groups and Steering Rules.

1. Select **Objects** > **Address Groups** and Set Up Dynamic Address Groups on Panorama for each tenant's virtual machines. For example, this use case has two security groups per tenant; one security group for the web servers and the other security group for the application servers.
2. Select **Policies** > **Security** > **Pre Rules** to set up security policy rules for sending traffic to the VM-Series firewall.
3. Select **Panorama** > **VMware** > **NSX-V** > **Steering Rules** and click **Auto-Generate Steering Rules**.
4. **Commit** your changes

STEP 5 | Prepare the ESXi Host for the VM-Series Firewall.

The ESXi hosts in the cluster must have the necessary NSX-V components that allow the NSX-V firewall and the VM-Series firewall to work together. The NSX-V Manager will install the components— the Ethernet Adapter Module (.eam) and the SDK —required to deploy the VM-Series firewall.

STEP 6 | Deploy the Palo Alto Networks NGFW Service.

1. Select **Networking and Security** > **Installation** > **Service Deployments**.
2. Click **New Service Deployment** (green plus icon), and select the service definition for the Palo Alto Networks next generation firewall you want to deploy, **Palo Alto Networks**

 ©2024 Palo Alto Networks, Inc.

NGFW Test 1 in this example, make your selections including the appropriate ESXi cluster to which you want to deploy the firewall and click **Finish**.

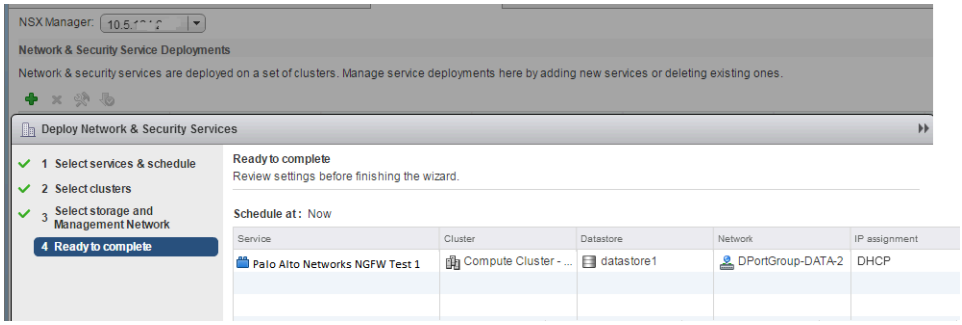

3. Verify that the NSX-V Manager reports the **Installation Status** as **Successful**.
4. Verify that the VM-Series firewall is successfully deployed.
   1. On the vCenter server, select **Hosts and Clusters** to check that every host in the cluster(s) has one instance of the firewall.
   2. View the management IP address(es) and the PAN-OS version running on the firewall directly from vCenter server. VMware Tools is bundled with the PAN-OS software image and is automatically enabled when you launch the VM-Series firewall.

STEP 7 | Apply Security Policies to the VM-Series Firewall.

1. Create Dynamic Address groups for each tenant on Panorama. The dynamic address group(s) that match on the name of the security group(s) you defined on the NSX-V Manager.
   1. On Panorama, select **Objects** > **Address Groups**.
   2. Select the correct **Device Group** from the drop-down and click **Add**.
   3. Add a **Name** for the address group and set Type as **Dynamic** and **Add Match Criteria**. Verify that you select the correct tags for each tenant, the tag includes the service

 ©2024 Palo Alto Networks, Inc.

profile ID, the security group name and the security group ID. For example, for this use case there are four dynamic address groups:

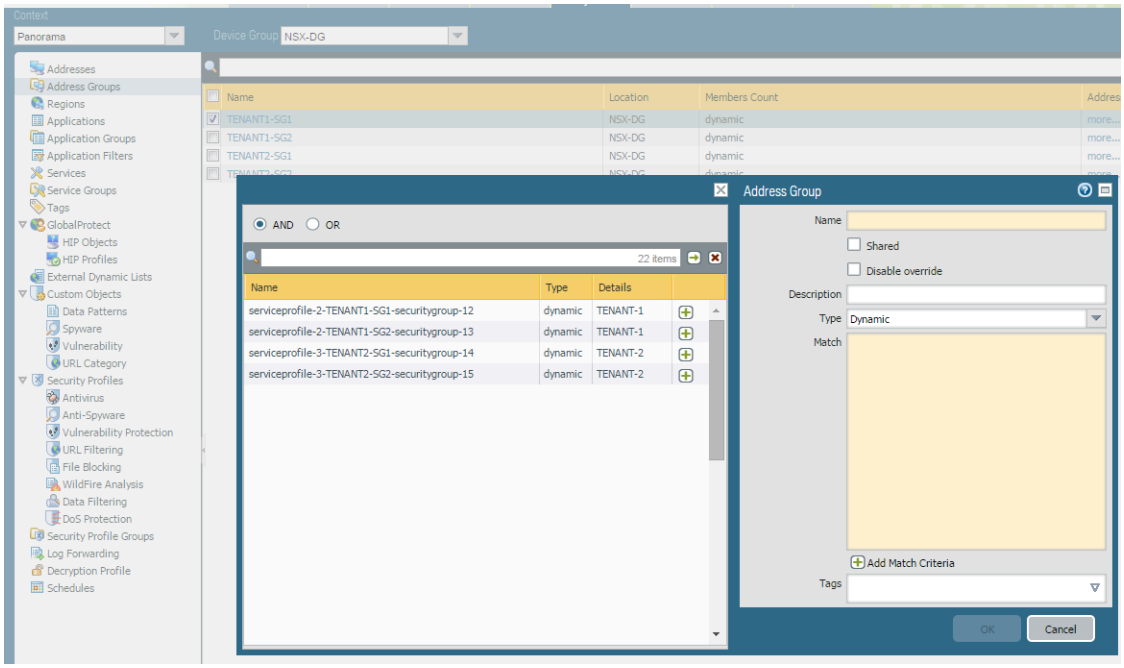

2. On Panorama, create security policy rules and use the dynamic address groups as source or destination address objects in security policy rules and push it to the firewalls.

   1. Select **Policies** > **Security** > **Prerules** and click **Add**.
   2. Create rules for each tenant. This use case has the following policy rules:

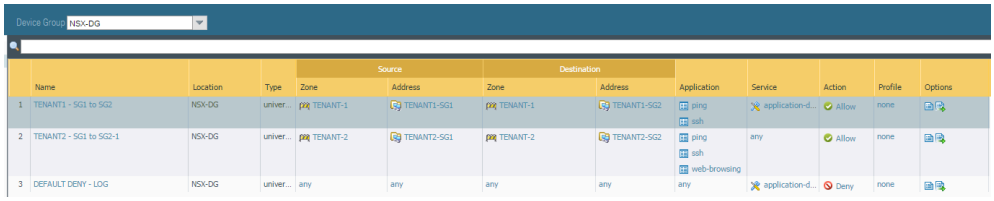

3. Click Commit, and select Commit Type as Device Groups. Select the device group, NSX-DG in this example and click OK.

STEP 8 | Verify that traffic from each tenant is secured.

1. Log in to the CLI on the firewall and enter the following command to view the subinterfaces on the firewall:

```
show interface all
total configured hardware interfaces: 2
name id    speed/duplex/state       mac address
--------------------------------------------------
ethernet1/1             16    auto/auto/up d4:f4:be:c6:af:10
ethernet1/2             17    auto/auto/up d4:f4:be:c6:af:11
aggregation groups: 0
total configured logical interfaces: 6

name                id vsys zone         forwarding
```

 ©2024 Palo Alto Networks, Inc.

```
------------------- ----- ---- -----------------
ethernet1/1       16    1                 vwire:ethernet1/2
ethernet1/1.3     4099  1    TENANT-1   vwire:ethernet1/2.3
ethernet1/1.4      4100   1   TENANT-2     vwire:ethernet1/2.4

ethernet1/2         17    1 vwire:ethernet1/1
ethernet1/2.3     4355  1    TENANT-1 vwire:ethernet1/1.3
ethernet1/2.4       4356  1 TENANT-2 vwire:ethernet1/1.4
```

2. On the web interface of the VM-Series firewall, select **Objects** > **Address Groups** and verify that you can view the IP address for the members of each Dynamic Address Group. The following is an example of duplicate IP addresses in dynamic address groups across both tenants.

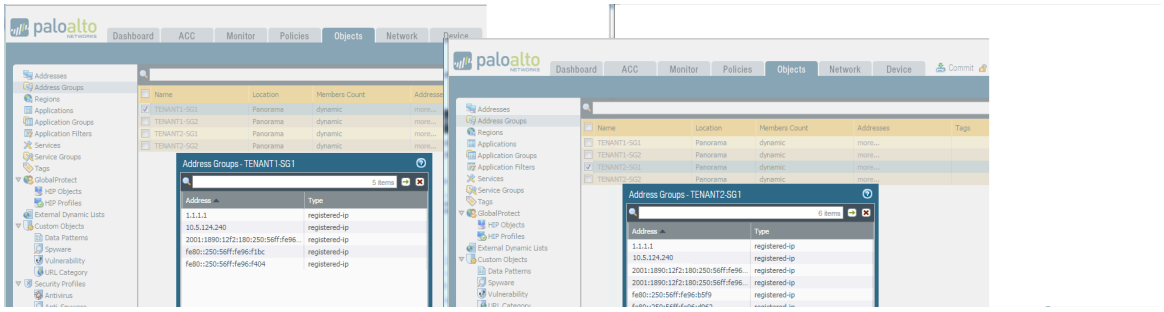

3. View the **ACC** and the **Monitor** > **Logs** > **Traffic**. Filter on the zone name to ensure that traffic from the virtual machines for each tenant is secured.

Use Case: Shared Security Policies on Dedicated Compute Infrastructure

If you are a Managed Service Provider who needs to secure a large enterprise (*tenant*) with multiple departments (*sub-tenants*), and each tenant requires dedicated compute infrastructure and security policy rules, you need to create a service definition for each tenant.

In this use case, each tenant—Oak and Maple— has a dedicated ESXi cluster. And each tenant has sub-tenants—Dev, QA, and Prod—whose workloads are deployed in the cluster. You need to define two service definitions to allow the VM-Series firewalls for each tenant to have Security policies for their respective ESXi clusters. The service definition for each tenant includes multiple zones (with corresponding virtual wire subinterface pairs) for isolating traffic from each sub-tenant. Each zone is mapped to a service profile on the NSX-V Manager, which allows the firewall to distinguish traffic from the virtual machines for each sub-tenant and to enforce zone-based security policy rules within the common set of policy rules for the tenant. Zone-based policies in combination with the Dynamic Address groups also allow you to secure sub-tenants who may have overlapping networks, and hence have duplicate IP addresses. To uniquely identify virtual machines assigned to each sub-tenant and successfully enforce policy, the NSX-V Manager provides the service profile and security group to which a virtual machine belongs as match criteria in dynamic address groups on Panorama. For more information, see Policy Enforcement using Dynamic Address Groups.

You can also configure role-based access control using access domains on Panorama. Access domains allow you to control administrative access to specific device groups (to manage policies and objects) and template stacks (to manage network and device settings), so that each tenant

 ©2024 Palo Alto Networks, Inc.

administrator can manage the configuration for their VM-Series firewalls. Role-based access also allows you to limit log visibility for the respective tenant only.

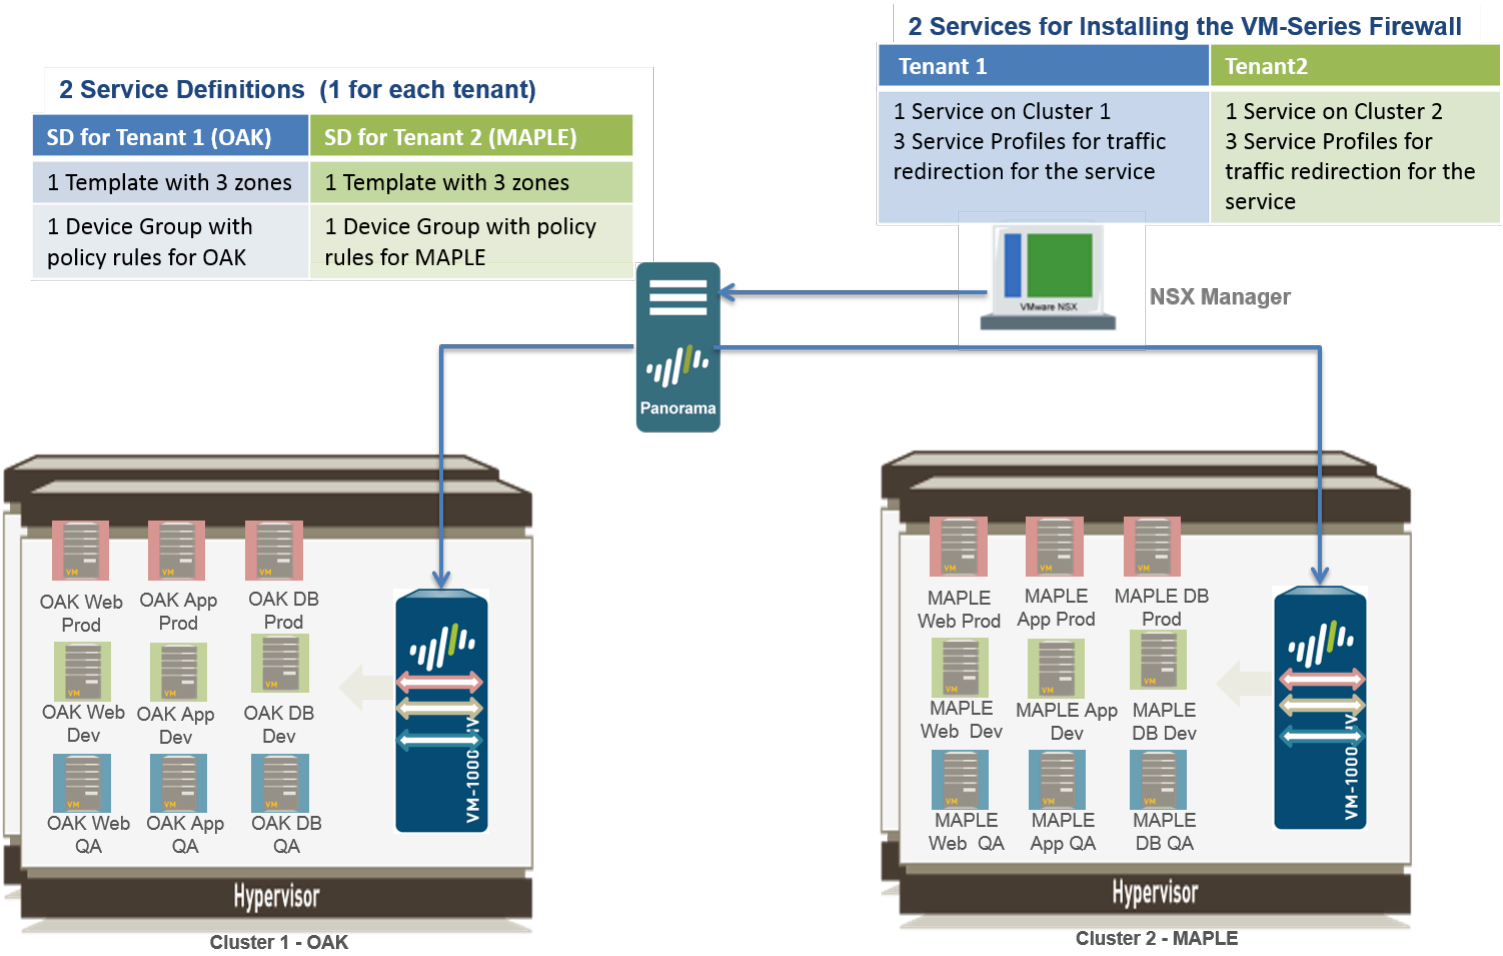

STEP 1 | Enable Communication Between the NSX-V Manager and Panorama.

This is one-time task and is required if you have not enabled access between the NSX-V Manager and Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Create Template(s) and Device Group(s) on Panorama.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Templates** to add template stacks. This use case has two template stacks named NSX-Template-MAPLE and NSX-Template-OAK.
3. Select **Panorama** > **Device Groups** and add device groups. This use case has two device groups named NSX-DG-OAK and NSX-DG-MAPLE.
4. Create NSX-V service profile zones within each template stack. To isolate traffic for each tenant in this use case, you need three zones for each tenant.
   1. Select **Network** > **Zones**.
   2. Select the correct template stack in the **Template** drop-down.
   3. Select **Add** and enter a zone **Name**. For example, **Tenant1**.
   4. Sets the interface **Type** to **Virtual Wire**.
   5. Click **OK**.
   6. Repeat the steps a-e to add additional zones for each sub-tenant.
   7. Verify that the zones are attached to the correct template stack.

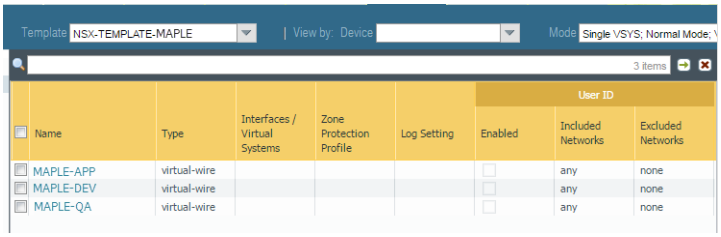

5. Create a service profile zone for each other template stack.

STEP 3 | Create the Service Definitions on Panorama.

1. Select **Panorama** > **VMware** > **NSX-V** > **Service Definitions**.
2. Select **Add**. Fill in the details for the service definition for each tenant. In this example, the two service definitions are Palo Alto Networks - Maple and Palo Alto Networks - Oak.

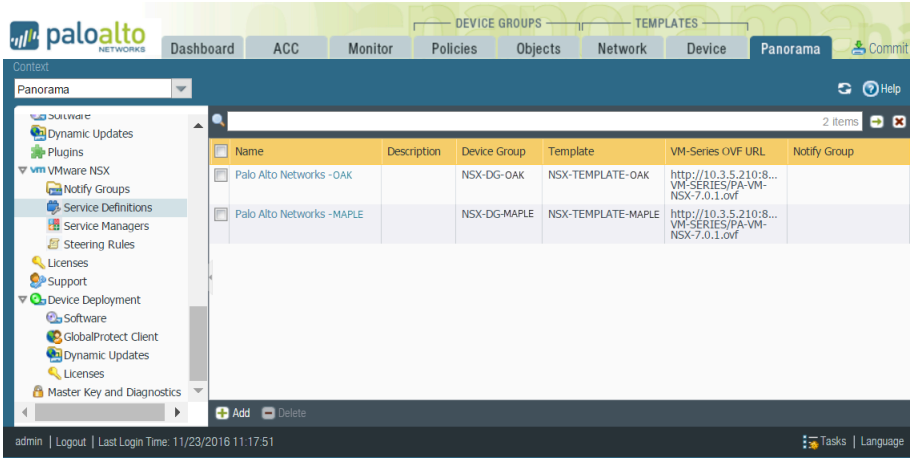

3. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Create Security Groups and Steering Rules.

1. Select **Objects** > **Address Groups** and Set Up Dynamic Address Groups on Panorama for each tenant's virtual machines. For example, this use case has two security groups per tenant; one security group for the web servers and the other security group for the application servers.

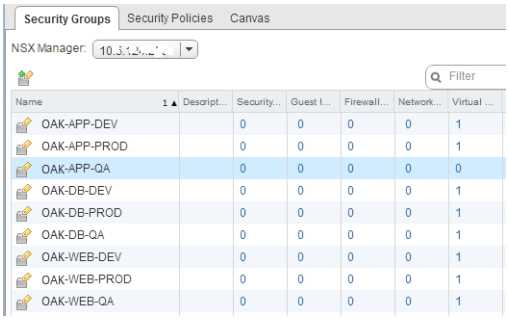

2. Select **Policies** > **Security** > **Pre Rules** to set up security policy rules for sending traffic to the VM-Series firewall.
3. Select **Panorama** > **VMware** > **NSX-V** > **Steering Rules** and click **Auto-Generate Steering Rules**.
4. **Commit** your changes

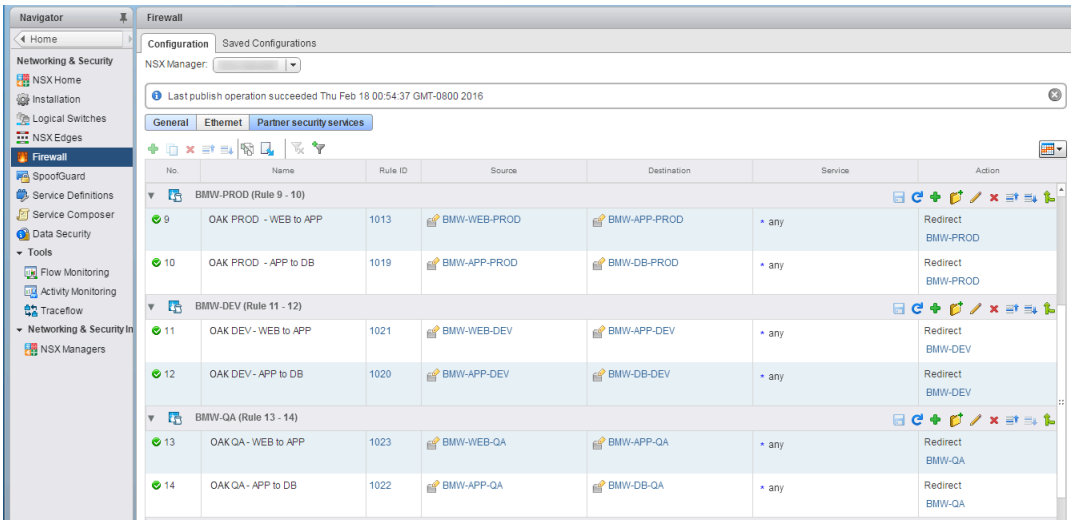

STEP 5 | Prepare the ESXi Host for the VM-Series Firewall

The ESXi hosts in the cluster must have the necessary NSX-V components that allow the NSX-V firewall and the VM-Series firewall to work together. The NSX-V Manager will install the components— the Ethernet Adapter Module (.eam) and the SDK —required to deploy the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Deploy the Palo Alto Networks NGFW Service

1. Select **Networking and Security** > **Installation** > **Service Deployments**.
2. Click **New Service Deployment** (green plus icon), and select the service definition for the Palo Alto Networks next generation firewall you want to deploy, **Palo Alto Networks NGFW Test 1** in this example, make your selections and click **Finish**.
3. Verify that the NSX-V Manager reports the **Installation Status** as **Successful**.

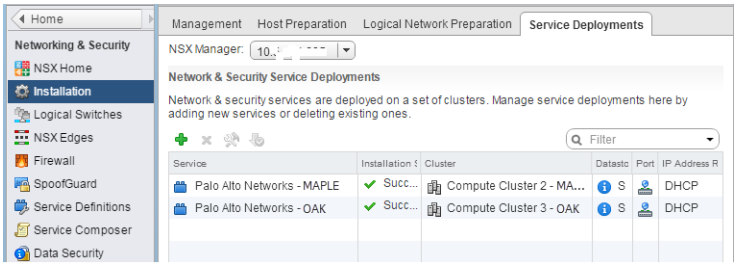

4. Verify that the VM-Series firewall is successfully deployed.
   1. On the vCenter server, select **Hosts and Clusters** to check that every host in each cluster has one instance of the firewall.
   2. View the management IP address(es) and the PAN-OS version running on the firewall directly from vCenter server. VMware Tools is bundled with the PAN-OS software image and is automatically enabled when you launch the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Apply Security Policies to the VM-Series Firewall

1. Create dynamic address groups for each sub-tenant on Panorama. The dynamic address group(s) match on the name of the security group(s) you defined on the NSX-V Manager.
   1. On Panorama, select **Objects** > **Address Groups**.
   2. Select a **Device Group** from the drop-down and click **Add**.
   3. Add a **Name** for the address group and set Type as **Dynamic** and **Add Match Criteria**. For ease of managing these groups, use the same name for the dynamic address group as that of the security group on the NSX-V Manager.

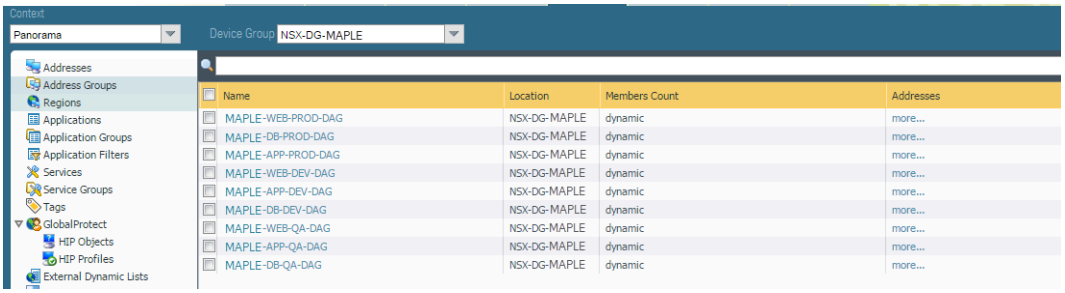

   4. Create the dynamic address groups for the sub-tenants for the other tenant, Oak in this example.
2. On Panorama, create Security policies and use the dynamic address groups as source or destination address objects in security policy rules and push it to the firewalls.
   1. Select **Policies** > **Security** > **Pre Rules**.
   2. Select a **Device Group** from the drop-down and click **Add**.
   3. Create rules for each sub-tenant. Make sure to keep the source and destination zone the same in a policy rule. To ensure that only the application that is running on the server is allowed, allow the service on the application-default port only.

      This use case has the following policy rules for the tenant Maple:

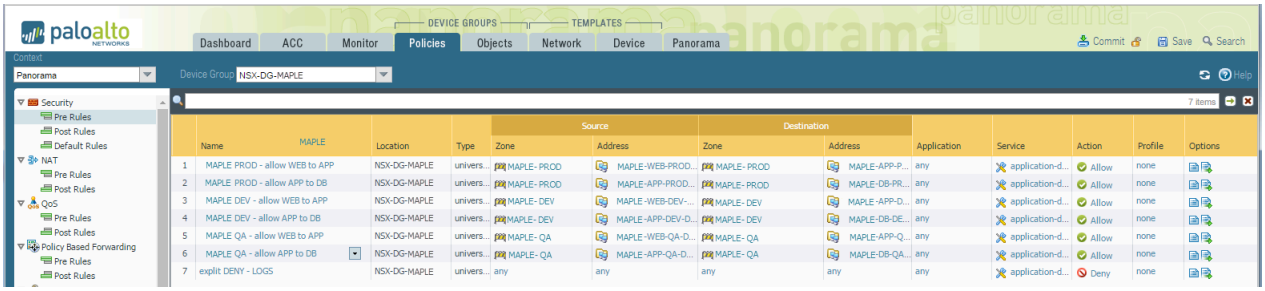

3. Select the other **Device Group** from the drop-down and create the Security policies for the each sub-tenant for the other tenant, Oak in this example.
4. Click Commit, and select Commit Type as Device Groups. Select the device groups, NSX-DG-OAK and NSX-DG-MAPLE in this example and click OK.

   The commit pushes the Security policies to the firewalls that belong to each device group, and they can enforce policy on the traffic redirected by the NSX-V Manager.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Verify that traffic from each tenant is secured.

1. On Panorama, go to **Monitor** > **Logs** > **Traffic** and **Monitor** > **Logs** > **Threat** to view the Traffic logs and Threat logs. Select the device group for a tenant and sort on the Zone name for full visibility in to traffic from each sub-tenant.
2. On Panorama, use the ACC for visibility into traffic patterns and actionable information on threats. Use the widgets and filters to interact with the data on the ACC.
3. On the VM-Series firewall, select **Objects** > **Address Groups** to view the IP address for the members of each Dynamic Address Group.

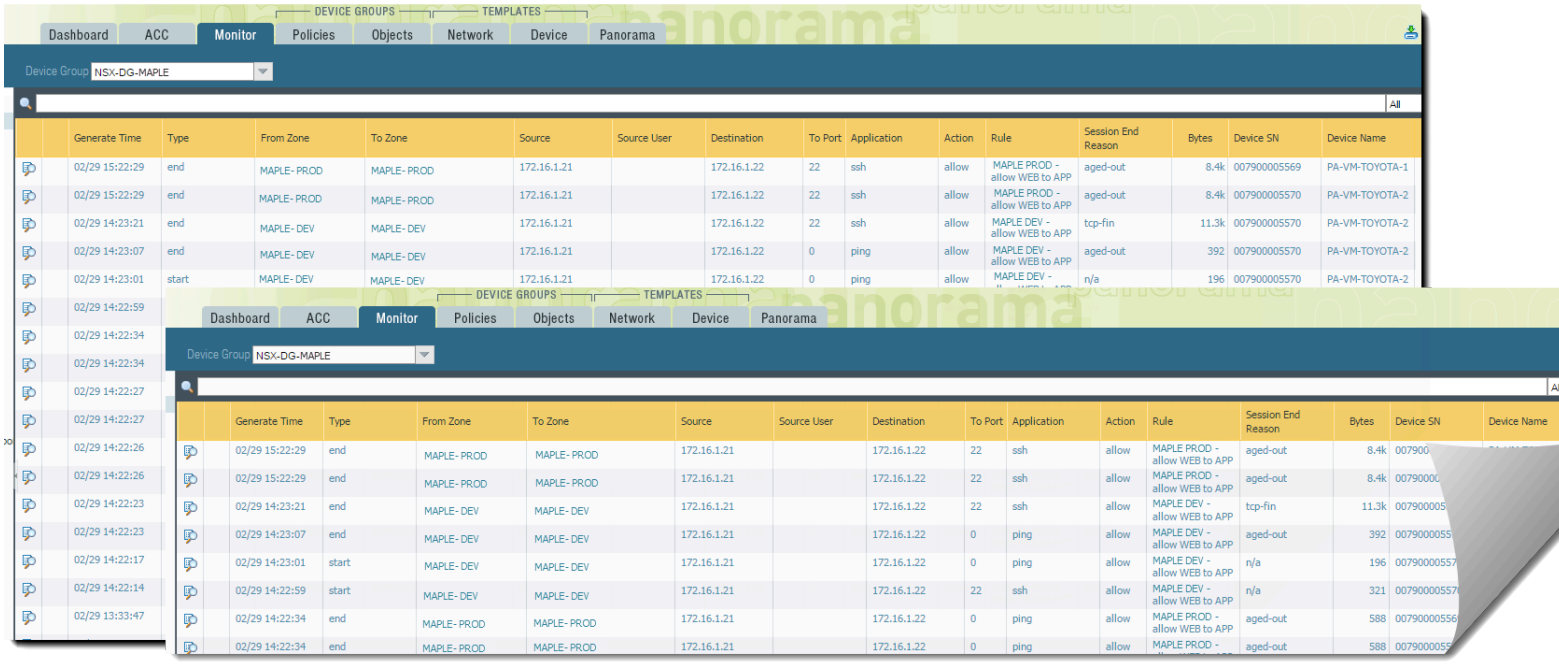

STEP 9 | (Optional) Enable role-based access for tenant administrators to manage the configuration and policies for the VM-Series firewalls.

1. Create an access domain. An access domain allows you to restrict admin access to a specific device group and template stack. In this example, you create two access domains and restrict access to the device group and template stack for the respective tenant.
2. Configure an admin role for **Device Group and Template** role and allow the administrator to manage the access domain. The administrator can only manage the firewalls that belong to the access domain.
3. Create an administrative account and associate the access domain and admin role with the account.

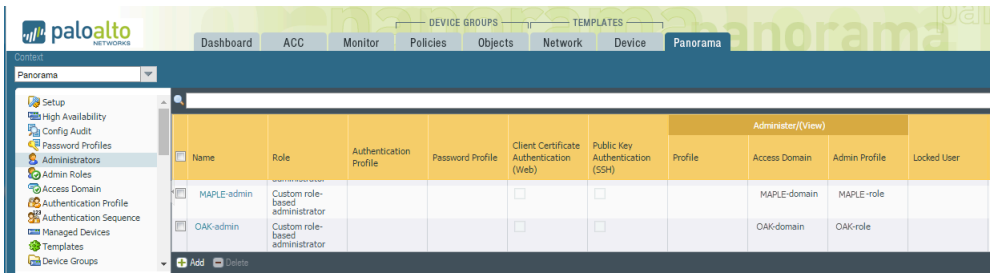

 ©2024 Palo Alto Networks, Inc.

Dynamic Address Groups—Information Relay from NSX-V Manager to Panorama

To enforce security policies in a VM-Series and NSX-V integrated data center, Panorama must be able to obtain information on the changes in the virtual landscape. As new virtual machines are deployed, changed, or deleted, the NSX-V Manager informs Panorama of IP addresses added, removed from security groups on the NSX-V Manager. Panorama in turn then, pushes this information to the VM-Series firewalls. Dynamic address groups referenced in firewall policies match against this information to determine the members that belong to the group. This process allows the firewall to enforce context-aware security policy, which secures traffic to and from these virtual machines. For details on dynamic address groups, see Policy Enforcement using Dynamic Address Groups.

The following diagram illustrates how the information is relayed from the NSX-V Manager to Panorama.

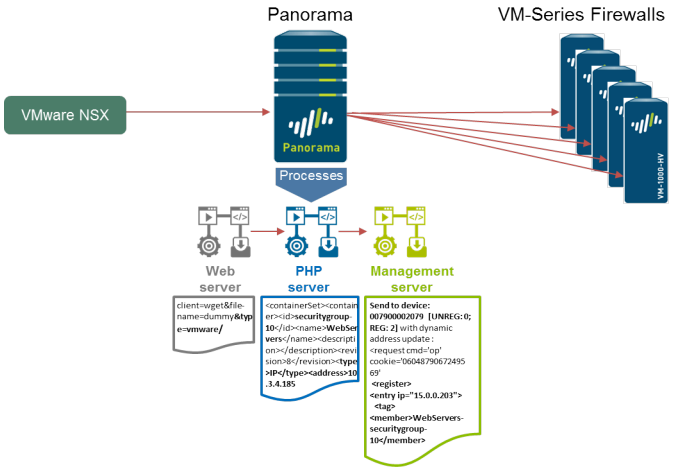

To understand this process, let's trace the information update sent from the NSX-V Manager to Panorama when a new server is added to a security group. Use the elements highlighted within the output in each phase of this example, to troubleshoot where the process failed.

STEP 1 | To view the updates in real-time, log in to the Panorama CLI.

Log in to the Command Line Interface on Panorama.

STEP 2 | Verify that the request from the NSX-V Manager is routed to the web server on Panorama.

To check the webserver-log on Panorama during an NSX-V Security Group update, use the following command:

```
admin@Panorama> tail follow yes webserver-log cmsaccess.log
127.0.0.1 - - [Wed Dec 03 14:24:11 2014 PST] "POST /unauth/php/
RestApiAuthenticator.php HTTP/1.1" 200 433
```

 ©2024 Palo Alto Networks, Inc.

```
127.0.0.1 - - [Wed Dec 03 14:24:11 2014 PST] "PUT /api/
index.php?client=wget&file-name=dummy&type=vmware/vmware/2.0/si/
serviceprofile/serviceprofile-1/containerset HTTP/1.0" 200 446
```

If your output does not include the elements above, check for routing issues. Ping the Panorama from the NSX-V Manager and check for ACLs or other network security devices that might be blocking the communication between the NSX-V Manager and Panorama.

STEP 3 | Verify that the request is parsed by the PHP daemon on Panorama.

1. Enable debug using the following URL: **`https://<Panorama_IP>/php/utils/debug.php`**

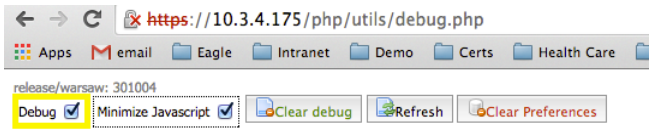

2. From the CLI, enter the following command to view the logs generated by the PHP server:

```
admin@Panorama> tail follow yes mp-log php.debug.log
[2014/12/03 14:24:11]
<request cmd="op" cookie="0604879067249569" refresh="no">
  <operations xml="yes">
    <show>
      <cli>
...
   <request>
      <partner>
        <vmware-service-manager>
          <update>
            <method>PUT</method>
            <type>update</type>
            <username>_vsm_admin</username>
           <password>4006474760514053</password>
<url>/vmware/2.0/si/serviceprofile/serviceprofile-1/
containerset</url>
            <data><![CDATA[
<containerSet><container><id>securitygroup-10</id><name>WebServers</
name><description></description><revision>8</revision><type>IP</
type><address>10.3.4.185</address><address>10.3.4.186</
address><address>15.0.0.203</address><address>15.0.0.202</address></
container></containerSet>]]></data>
```

```
</update>
        </vmware-service-manager>
      </partner>
    </request>
  </operations>
</request>
```

 ©2024 Palo Alto Networks, Inc.

STEP 4 | The information is processed by the Management server on Panorama.

1. Enable debugging on the management server using the following command:

```
admin@Panorama> debug management-server on debug
```

2. Enter the following command to view the logs generated by the configd log:

```
admin@Panorama> tail follow yes mp-log configd.log
```

3. In the output check that the update was relayed from the PHP daemon to the management server daemon.

```
2014-12-03 14:24:11.143 -0800 debug:
 pan_job_progress_monitor(pan_job_mgr.c:3694): job-monitor:
 updated 0 jobs……2014-12-03 14:24:11.641 -0800 debug:
 recursive_add_params(pan_op_ctxt.c:158): > 'url'='/
vmware/2.0/si/serviceprofile/serviceprofile-1/containerset'
2014-12-03 14:24:11.641 -0800 debug:
 recursive_add_params(pan_op_ctxt.c:158): > 'data='
<containerSet><container><id>securitygroup-10</id><name>WebServers</
name><description></description><revision>8</revision><type>IP</
type><address>10.3.4.185</address><address>10.3.4.186</
address><address>15.0.0.203</address><address>15.0.0.202</address></
container></containerSet>'
2014-12-03 14:24:11.641 -0800 Received vshield update: PUT /
vmware/2.0/si/serviceprofile/serviceprofile-1/containerset
Received dynamic address update from VSM:
<request cmd='op' cookie='0604879067249569'
 client="xmlapi"><operations xml='yes'><request>
    <partner>
        <vmware-service-manager>
            <update>
                <method>PUT</method>
        <type>update</type>
<username>_vsm_admin</username>
<password>4006474760514053</password><url>/vmware/2.0/si/
serviceprofile/serviceprofile-1/containerset</url><data><!
[CDATA[
<containerSet><container><id>securitygroup-10</id><name>WebServers</
name><description></description><revision>8</revision><type>IP</
type><address>10.3.4.185</address><address>10.3.4.186</
address><address>15.0.0.203</address><address>15.0.0.202</address></
container></containerSet>]]>
```

```
</data></update>
```

4. Look for the list of IP addresses and security group tags.

```
2014-12-03 14:24:11.646 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3721): ip:
10.3.4.185
```

 ©2024 Palo Alto Networks, Inc.

```
2014-12-03 14:24:11.646 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3738): tag:
WebServers-securitygroup-10
2014-12-03 14:24:11.646 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3721): ip:
15.0.0.202
2014-12-03 14:24:11.646 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3738): tag:
WebServers-securitygroup-10
pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3738): tag:
DomainControllers-securitygroup-16
2014-12-03 14:24:11.647 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3721): ip:
15.0.0.201
2014-12-03 14:24:11.648 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3738): tag:
SQLServers-securitygroup-11
2014-12-03 14:24:11.665 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3738): tag:
SharePointServers-securitygroup-13
2014-12-03 14:24:11.665 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3721): ip:
10.3.4.187
2014-12-03 14:24:11.665 -0800 debug:
 pan_cfg_mongo_sel_ip_taglist_by_tag_rev(src_cms/
pan_cfg_mongo_tables.c:3738): tag:
SharePointServers-securitygroup-13
...
```

5. Finally, verify that the update was relayed from the management server daemon to the managed firewalls.

```
Send to device: 007900002079 [UNREG: 0; REG: 2] with dynamic
 address update : <request cmd='op' cookie='0604879067249569'
 target-…. <register>
<entry ip="15.0.0.203">
  <tag>
<member>WebServers-securitygroup-10</member>
  </tag>
</entry>
<entry ip="10.3.4.186">
  <tag>
<member>WebServers-securitygroup-10</member>
  </tag>
</entry>
</register>
```

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on VMware NSX-T (North-South)

The VM-Series firewall on VMware NSX-T integrates the Palo Alto next-generation firewalls and Panorama with ESXi host servers to provide comprehensive visibility and safe application enablement of all north-south traffic in your NSX-T software-defined datacenter.

The following topics provide information about the VM-Series firewall on VMware NSX-T:

- Supported Deployments of the VM-Series Firewall on VMware NSX-T (North-South)
- Components of the VM-Series Firewall on NSX-T (North-South)
- Deploy the VM-Series Firewall on NSX-T (North-South)
- Extend Security Policy from NSX-V to NSX-T

Supported Deployments of the VM-Series Firewall on VMware NSX-T (North-South)

You can deploy one or more instances of the VM-Series firewall as a partner service in your VMware NSX-T Data Center. Attach a VM-Series firewall to any tier-0 or tier-1 logical router to protect north-south traffic. You can deploy the VM-Series firewall as standalone service instance or two firewalls in a high-availability (HA) pair. Panorama manages the connection with NSX-T Manager and the VM-Series firewalls deployed in your NSX-T software-defined datacenter.

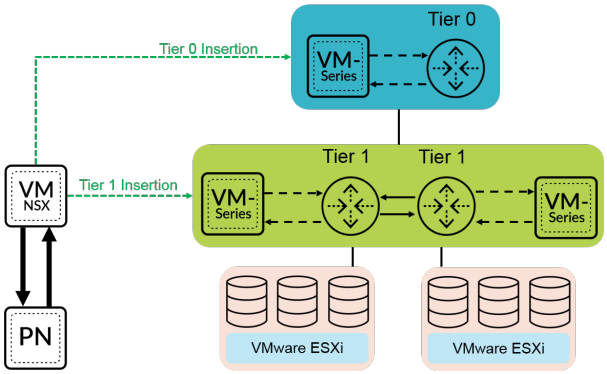

- Tier-0 Insertion—Tier-0 insertion deploys a VM-Series firewall to a tier-0 logical router, which processes traffic between logical and physical networks. When you deploy the VM-Series firewall with tier-0 insertion, NSX-T Manager uses the deployment information you configured on Panorama to attach a firewall to a tier-0 logical router in virtual wire mode.
- Tier-1 Insertion—Tier-1 insertion deploys a VM-Series firewall to a tier-1 logical router, which provides downlink connections to segments and uplink connection to tier-0 logical routers. NSX-T Manager attaches VM-Series firewalls deployed with tier-1 insertions to a tier-1 logical router in virtual wire mode.

After deploying the firewall, you configure traffic redirection rules that send traffic to the VM-Series firewall when crossing a tier-0 or tier-1 router. Security policy rules that you configure on Panorama are pushed to managed VM-Series firewalls and then applied to traffic passing through the firewall.

 ©2024 Palo Alto Networks, Inc.

Components of the VM-Series Firewall on NSX-T (North-South)

The following tables show the components of this joint Palo Alto Networks and VMware NSX-T solution.

VMware Components	
vCenter/ESXi	The vCenter server is the centralized management tool for the vSphere suite. ESXi is a hypervisor that enables compute virtualization. Refer to VMware's Compatibility Matrix for vCenter compatibility with your version of NSX-T.
NSX-T Manager	VMware NSX-T Data Center 2.4.0 and later must be installed and registered with the vCenter server. The NSX-T Manager is required to deploy the VM-Series firewall on the ESXi hosts within a ESXi cluster.

Palo Alto Networks Components	
PAN-OS	The VM-Series base image (PA-VM-NST-9.1.zip) is required for deploying the VM-Series firewall on NSX-T. The minimum system requirement for deploying the VM-Series firewall for NSX on the ESXi server depends on your VM-Series model. See VM-Series Models for the minimum hardware requirements for your VM-Series model.
Panorama Panorama must be running the same release version or later version that the firewalls that it will manage.	The VM-Series firewall on NSX-T requires Panorama 9.1 or later. Panorama is the centralized management tool for the Palo Alto Networks next-generation firewalls. In this solution, Panorama works with the NSX-T Manager to deploy, license, and centrally administer—configuration and policies—the VM-Series firewall for NSX-T. Panorama must be able to connect to the NSX-T Manager, the VM-Series firewalls and the Palo Alto Networks update server. See the Panorama Administrator's Guide for information about deploying your Panorama appliance.

 ©2024 Palo Alto Networks, Inc.

Palo Alto Networks Components	
Panorama Plugin for VMware NSX	3.0.0 or later
VM-Series Plugin	1.0.6 or later
VM-Series Firewall Models	The VM-100, VM-300, VM-500, and VM-700 support NSX-T.

Deploy the VM-Series Firewall on NSX-T (North-South)

Complete the following tasks to secure North-South traffic in your NSX-T environment with the VM-Series firewall.

The following procedure refers to NSX-T Manager 3.0

.

- Install the Panorama Plugin for VMware NSX
- Enable Communication Between NSX-T Manager and Panorama
- Create Template Stacks and Device Groups on Panorama
- Configure the Service Definition on Panorama
- Deploy the VM-Series Firewall
- Direct Traffic to the VM-Series Firewall
- Apply Security Policy to the VM-Series Firewall on NSX-T
- Use vMotion to Move the VM-Series Firewall Between Hosts

Install the Panorama Plugin for VMware NSX

Download and install the Panorama Plugin for VMware NSX. See the Compatibility Matrix before installing or upgrading your plugin.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panorama HA peers, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away

 ©2024 Palo Alto Networks, Inc.

- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  ```
  request plugins dau plugin-name <plugin-name> unblock-device-push yes
  ```

  You can cancel this command by executing:

  ```
  request plugins dau plugin-name <plugin-name> unblock-device-push no
  ```

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Select **Panorama** > **Plugins**. See the Compatibility Matrix before installing or upgrading your plugin.

STEP 2 | Select **Check Now** to retrieve a list of available updates.

STEP 3 | Select **Download** in the Action column to download the plugin.

STEP 4 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

Enable Communication Between NSX-T Manager and Panorama

Complete the following procedure to enable communication between Panorama and NSX-T Manager. You can connect your Panorama to up to 16 NSX-T Managers. If you are connecting your Panorama to multiple NSX-T Managers, you must carefully plan your device group hierarchy and template stacks and consider how they interact with the other components needed for deployment. Service definitions reference device groups and template stacks and push that information to the firewalls in the related ESXi clusters.

STEP 1 | (Optional) Bypass proxy server settings, configured on Panorama under **Panorama** > **Setup** > **Services** > **Proxy Server**, for communication between Panorama and NSX-T Manager. This command allows Panorama to communicate directly with NSX-T Manager while maintaining proxied communication for other services.

1. Log in to the Panorama CLI.
2. Execute the following command to enable or disable proxy bypass.

   ```
   admin@Panorama> request plugins vmware_nsx proxy bypass {yes | no}
   ```

   Select **yes** to enable proxy bypass and **no** to disable proxy bypass. This is set to **no** by default.

STEP 2 | Log in to the Panorama web interface.

Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<*IP address*>).

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Set up access to the NSX-T Manager. Repeat this procedure for each NSX-T Manager to which you will connect Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers** and click **Add**.
2. Enter a descriptive **Name** for your NSX-T Manager.
3. (Optional) Add a **Description** for NSX-T Manager.
4. Enter the **NSX Manager URL**—NSX-T Manager cluster virtual IP address or FQDN—at which to access the NSX-T Manager.
5. Enter the **NSX Manager Login** credentials—username and password, so that Panorama can authenticate to the NSX-T Manager.
6. Click **OK**.

If you change your NSX-T Manager login password, ensure that you update the password on Panorama immediately. An incorrect password breaks the connection between Panorama and NSX-T Manager.

STEP 4 | Commit your changes to Panorama.

Select **Commit** and **Commit to Panorama**.

STEP 5 | Verify the connection status on Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers**.
2. Verify the message in the **Status** column.

   When the connection is successful, the status displays as **Registered**. This indicates that Panorama and the NSX-T Manager are in sync.

   The unsuccessful status messages are:

   - **No connection**: Unable to reach/establish a network connection to the NSX-T Manager.
   - **Invalid Credentials**: The access credentials (username and/or password) are incorrect.
   - **Out of sync**: The configuration settings defined on Panorama are different from what is defined on the NSX-T Manager.Click the link for details on the reasons for failure. For example, NSX-T Manager may have a service definition with the same name as defined on Panorama. To fix the error, use the service definition name listed in the error message to validate the service definition on the NSX-T Manager. Until the configuration on Panorama and the NSX-T Manager is synchronized, you cannot add a new service definition on Panorama.
   - **Connection Disabled**: The connection between Panorama and the NSX-T Manager was manually disabled.

Create Template Stacks and Device Groups on Panorama

To manage the VM-Series firewalls on NSX-T using Panorama, the firewalls must belong to a device group and a template stack. Device groups allow you to assemble firewalls that need similar policies and objects as a logical unit; the configuration is defined using the **Objects** and **Policies** tabs on Panorama. Use template stacks to configure the settings that are required for the VM-Series firewalls to operate on the network; the configuration is defined using the **Device**

 ©2024 Palo Alto Networks, Inc.

and **Network** tabs on Panorama. Each template stack used in your NSX-T configuration must be associated with a service definition.

Firewalls deployed in NSX-T have two default zones and two interfaces configured in virtual-wire mode. Ethernet1/1 is part of zone **south** and ethernet1/2 is part of zone **north**. To push policy rules from Panorama to managed firewalls, you must configure zones and interfaces matching those on the firewall in the corresponding template stack on Panorama.

STEP 1 | Add a device group or a device group hierarchy.

1. Select **Panorama** > **Device Groups**, and click **Add**. You can also create a device group hierarchy.
2. Enter a unique **Name** and a **Description** to identify the device group.
3. Click **OK**.
4. Click **Commit** and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 2 | Add a template.

1. Select **Panorama** > **Templates**, and click **Add**.
2. Enter a unique **Name** and a **Description** to identify the template.
3. Click **OK**.
4. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 3 | Create a template stack.

1. Select **Panorama** > **Templates**, and click **Add Stack**.
2. Enter a unique **Name** and a **Description** to identify the template.
3. Click **Add** to add the template you created previously.
4. Click **OK**.
5. Click **Commit**, and select **Commit to Panorama** to save the changes to the running configuration on Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Configure the virtual wire, interfaces, and zones. Ensure that you select the correct template from the drop-down shown below. The objects you create must meet the following criteria:

If you change the default virtual wire or zone names, the virtual wire and zones on Panorama must match the names used on the firewall.

- Use **ethernet1/1** and **ethernet1/2**.
- The virtual wire object named **vw1**.
- The first zone named **south**, type **virtual-wire**, and contain **ethernet1/1**.
- The second zone named **north**, type **virtual-wire**, and contain **ethernet1/2**.

Repeat this process for each template in your deployment.

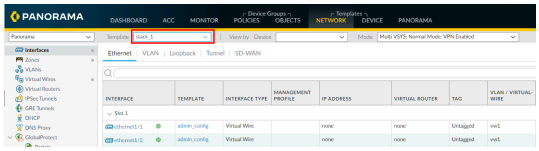

STEP 5 | Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

Configure the Service Definition on Panorama

A service definition specifies the configuration for the VM-Series firewalls installed in your NSX-T data center environment. The service definition must include the device group, a template stack, and an OVF URL.

STEP 1 | Add a new service definition.

You can create up to 32 service definitions on Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Definitions**.
2. Select **Add** to create a new service definition.
3. Enter a descriptive **Name** for your service definition.
4. (Optional) Add a **Description** that identifies the function or purpose for the VM-Series firewalls that will be deployed using this service definition.

STEP 2 | Assign a device group and a template stack to the service definition.

Make sure to Create Template Stacks and Device Groups on Panorama.

Because the firewalls deployed in this solution will be centrally administered from Panorama, you must specify the **Device Group** and the **Template Stack** that the firewalls belong to. All the

 ©2024 Palo Alto Networks, Inc.

firewalls that are deployed using this service definition belong to the specified template stack and device group.

1. Select the device group or device group hierarchy in the **Device Group** drop-down.
2. Select the template stack in the **Template** drop-down.

   You cannot reuse a template stack or a device group assigned to one service definition in another service definition.

STEP 3 | Specify the location of the OVF file.

Download the zip file, unzip it to extract and save the .ovf, mf and .vmdk files to the same directory. The ovf and vmdk files are used to deploy each instance of the firewall.

If needed, modify the security settings on the server so that you can download the file types. For example, on the IIS server modify the Mime Types configuration; on an Apache server edit the .htaccess file.

Do not change the Panorama service definition OVF path after a successful NSX Service Deployment of VM-Series firewalls. Changing the OVF path, after a successful VM-Series firewall deployment, can result in a NSX Service Deployment failed state. You may resolve this failure in NSX-T Manager, however this may cause all VM-Series firewalls to redeploy.

It is recommended that you use an OVF path name that scales and allows you to change the base image without impacting your deployed firewalls. Instead of a path such as **https://acme.com/software/PA-VM-NST.9.1.0.ovf**, use something such as **https://acme.com/software/PanoSvcDef1-Cluster1.ovf**. Using a static path reference will eliminate any future need to change the OVF path. It is recommended to create a path for each Panorama

 ©2024 Palo Alto Networks, Inc.

service definition (vSphere cluster) in your deployment and change the PAN-OS base images references on the web server as needed.

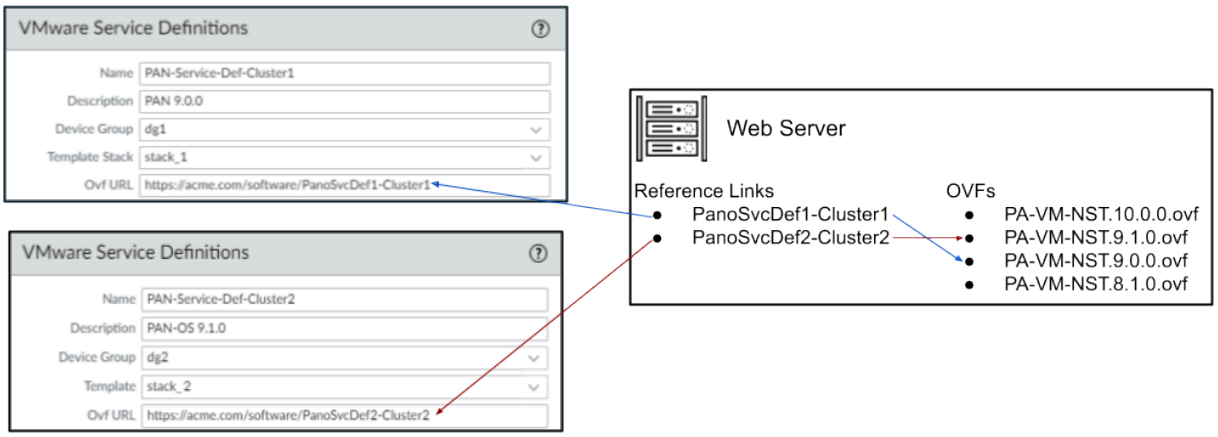

In **OVF URL**, add the location of the web server that hosts the ovf file. Both http and https are supported protocols.

Panorama must have network connectivity with the web server to retrieve the OVF file.

You can use the same ovf version or different versions across service definitions. Using different ovf versions across service definitions allows you to vary the PAN-OS version on the VM-Series firewalls in different ESXi clusters.

STEP 4 | Select **North South** as the **Insertion Type** for your firewall.

STEP 5 | To automatically retrieve a device certificate when the VM-Series firewall is deployed by NSX Manager, configure the device certificate.

Enable this option to apply a device certificate to newly deployed VM-Series firewalls. Only use this option when deploying the firewall using a base image OVF that supports device certificates. Panorama pushes the device certificate information to NSX Manager as part of the

 ©2024 Palo Alto Networks, Inc.

service definition. When a new firewall is deployed in NSX, the device certificate is installed on the firewall at bootup.

For list of OVFs that support device certificates for the VM-Series firewall on VMware NSX, see the Palo Alto Networks Compatibility Matrix.

If your OVF does not support a device certificate, disable this option.

1. If you have not done so already, log in to the Customer Support Portal and generate a Registration PIN and PIN ID.
2. Under **Device Certificate**, click **Enable**.
3. Copy the PIN ID and enter it into the **Device Certificate PIN ID** field.
4. Reenter the PIN ID into the **Confirm Device Certificate PIN ID** field.
5. Copy the PIN Value and enter it into the **Device Certificate PIN Value** field.
6. Reenter the PIN Value into the **Confirm Device Certificate PIN Value** field.

STEP 6 | Click **OK** to save the service definition.

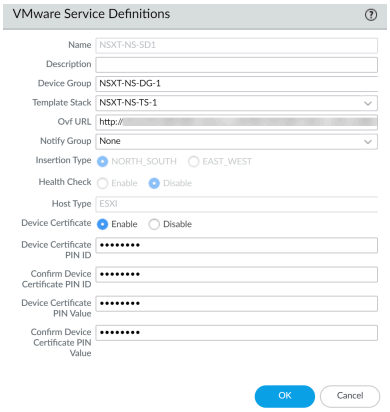

STEP 7 | Attach the service definition to the service manager.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Manager** and click the link of the service manager name.
2. Under Service Definitions, click **Add** and select your service definition from the drop-down.
3. Click **OK**.

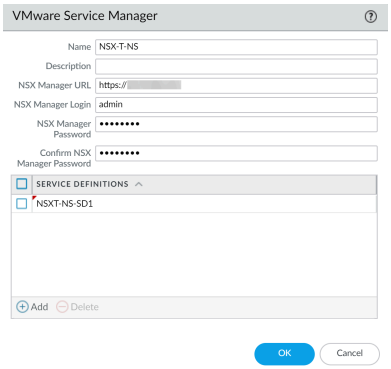

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Add the authorization code to license the firewalls.

1. Select **Panorama** > **Device Groups** and choose the device group you associated with the service definition you just created.
2. Under **Dynamically Added Device Properties**, add the authorization code you received with your order fulfillment email and, optionally, select None from the **SW Version** drop-down.

   When a new firewall is deployed on NSX-T it is automatically added to the device group, licensed using the authorization code you provided, and upgraded to the PAN-OS version you specified.

   On the support portal, you can view the total number of firewalls that you are authorized to deploy and the ratio of the number of licenses that have been used to the total number of licenses enabled by your authorization code.

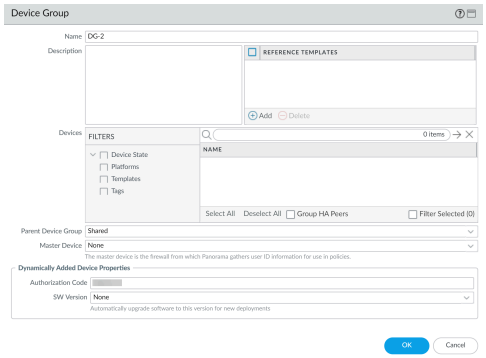

STEP 9 | **Commit to Panorama**.

STEP 10 | On the NSX-T Manager, verify that the service definition is available.

Select **System** > **Service Deployments** > **Catalog**. The service definition is listed as a Service Instance on the NSX-T Manager.

Deploy the VM-Series Firewall

After completing the configuration on Panorama, perform the following procedure to launch the VM-Series firewall in your NSX-T Data Center.

When deploying the VM-Series firewall on NSX-T in high availability, both firewalls are deployed to the same Device Group and Template Stack.

STEP 1 | Log in to NSX-T Manager.

STEP 2 | Select **System** > **Service Deployments** > **Deployment**.

STEP 3 | Select your service definition from the **Partner Service** drop-down.

STEP 4 | Click **Deploy Service**.

STEP 5 | Enter a descriptive **Service Deployment Name** for your VM-Series firewall.

STEP 6 | Select a tier-0 or tier-1 router under **Attachment Points**. NSX-T Manager attaches the VM-Series firewall to the selected router and redirects traffic passing through that router to

 ©2024 Palo Alto Networks, Inc.

the VM-Series firewall for inspection. You must select a router with no service insertion attached.

STEP 7 | Select a **Compute Manager**. The compute manager is the vCenter server managing your datacenter.

STEP 8 | Select a **Cluster**. You can deploy the VM-Series firewall on any cluster that does not include any Edge Transport Nodes.

STEP 9 | Select a **Datastore**.

STEP 10 | Configure your network settings.

1. Click **Edit Details** in the **Networks** column.
2. Select the **Primary Interface Network**.
3. Enter the **Primary Interface IP**.
4. Enter the **Primary Gateway Address**.
5. Enter the **Primary Subnet Mask**.
6. Click **Save**.

STEP 11 | NSX-T Manager prepopulates the **Deployment Specification** and **Deployment Template** based on the Partner Service you selected.

STEP 12 | Set the **Failure Policy** to Allow or Block. The failure policy defines how NSX-T Manager handles traffic that is directed to the VM-Series firewall if the firewall becomes unavailable.

STEP 13 | Select the **Deployment Mode** for your VM-Series firewall—Standalone or High Availability. If you have an edge node cluster and select High Availability, NSX-T Manager will deploy an additional VM-Series firewall on the standby edge node in addition to the firewall deployed on the active edge node.

STEP 14 | Click **Save** to deploy the VM-Series firewall.

STEP 15 | Verify that your firewalls connected to Panorama.

1. Log in to Panorama.
2. Select **Panorama** > **Managed Devices** > **Summary**.
3. Confirm that your firewalls are listed under the correct device group and the **Device State** shows **Connected**.

   The Device Name for the VM-Series firewall is displayed on Panorama as **PA-VM:<nsx.clusterid>** for NSX-T (N-S) deployment and as **PA-VM:<nsx.servicevmid>** for NSX-T (E-W) deployment.

STEP 16 | Set a secure password for the admin account on your VM-Series firewalls.

Each VM-Series firewall uses a default username and password (admin/admin), which is used for initial login. Upon logging in for the first time, you are prompted to set a new, more secure

 ©2024 Palo Alto Networks, Inc.

password. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

You can update the password on each firewall individually or all at once through Panorama.

- **Panorama**—on Panorama, you can change the default password for all firewalls in a template or delete the admin user and create a new username and password.
  1. Log in to Panorama
  2. Select **Device** > **Administrators** and select the **admin** user.
  3. **Delete** the user or click the user and enter a new password.
  4. If you changed the password, click **OK**.
  5. Select **Commit** > **Push to Devices** > **Edit Selections** > **Force Template Values**.
  6. Click **OK**.
- **Firewall**—this procedure must be repeated on each VM-Series firewall.
  1. Log in to the VM-Series firewall using the default username and password.
  2. Follow the prompts to reset the password.

Direct Traffic to the VM-Series Firewall

Complete the following procedure to direct traffic to your VM-Series firewall. For North-South traffic, redirection rules are stateless by default and cannot be changed. Additionally, NSX-T automatically creates a corresponding reflexive rule for return traffic.

When you deploy the VM-Series firewall for NSX-T North-South in HA mode, you must create a traffic redirection rule for both HA peers. Additionally, you must create the redirection rule for active peer first and the passive peer second.

The reflexive rule does not appear in the NSX-T web interface.

STEP 1 | Log in to NSX-T Manager.

STEP 2 | Verify that your are in **Policy** mode.

STEP 3 | Select **Security** > **North South Security** > **Network Introspection (N-S)**.

STEP 4 | Click **Add Policy**.

STEP 5 | Enter a descriptive **Name** for your policy.

STEP 6 | Select a VM-Series firewall service instance from the **Redirect To** drop-down. NSX-T Manager will automatically populate the **Applied To** field based on the service instance you select.

STEP 7 | Select your newly created policy.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Click **Add Rule**.

If your NSX-T environment has Edge Nodes in active-standby HA, you must create a redirect rule for each Edge Node. NSX-T does not automatically apply a redirect rule to the standby node in the event of a failover.

STEP 9 | Click on the **Name** field and enter a descriptive name for the rule.

STEP 10 | By default, the source is set to Any. Complete the following steps to specify a different source.

1. Click on the edit button in the **Source** column.
2. Select the group or groups to set as the Source or click Add Group to create a new group.
3. Click **Apply**.

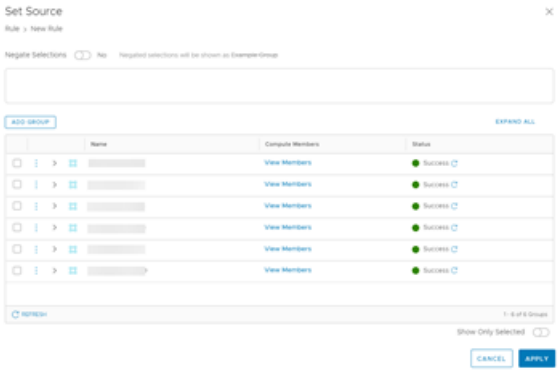

STEP 11 | By default, the destination is set to Any. Complete the following steps to specify a different destination.

1. Click on the edit button in the **Destination** column.
2. Select the group or groups to set as the Destination or click Add Group to create a new group.
3. Click **Apply**.

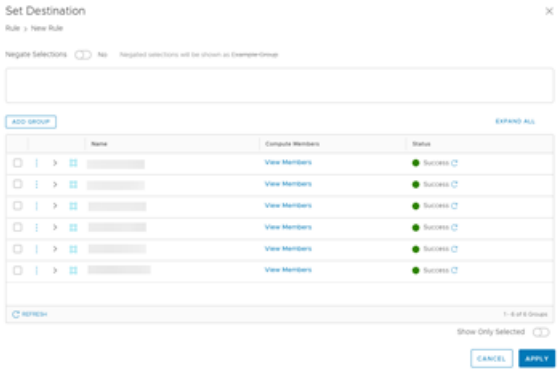

 ©2024 Palo Alto Networks, Inc.

STEP 12 | By default, Any service is redirected to the firewall. Complete the following steps to specify certain services and protocols.

1. Click on the edit button in the **Services** column.
2. Select the group or groups to set as the Service or click Add Service to create a new service.
3. Click **Apply**.

STEP 13 | Select **Redirect** from the **Action** drop-down to send traffic to your VM-Series firewall.

STEP 14 | **Enable** the rule. NSX-T Manager publishes the redirection rule you just created and automatically creates a reflexive rule for return traffic. The reflexive rule does not appear in the NSX-T Manager web interface.

STEP 15 | If your VM-Series firewalls are deployed in HA, create another rule for the passive HA peer.

If return traffic is not directed to the VM-Series firewall, manually configure a traffic redirection rule for return traffic.

Apply Security Policy to the VM-Series Firewall on NSX-T

Now that you have deployed the VM-Series firewall and created traffic redirection rules to send traffic to the firewall, you can use Panorama to centrally manage security policy rules on the VM-Series firewall.

STEP 1 | Log in to Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Create security policy rules.

By default, the firewall creates a rule that allows Bidirectional Forwarding Detection (BFD). Do not create a rule that blocks BFD. If BFD is blocked, NSX-T thinks that the firewall is unavailable.

1. Select **Policies** > **Security** > **Prerules**.
2. Select the **Device Group** that you created for managing the VM-Series firewalls on NSX-T in Create Template Stacks and Device Groups on Panorama.
3. Click **Add** and enter a **Name** and a **Description** for the rule. In this example, the security rule allows all traffic between the WebFrontEnd servers and the Application servers.
4. Select the **Source Zone** and **Destination Zone**.
5. For the **Source Address** and **Destination Address**, select or type in an address, static address group, or region.

   The VM-Series firewall on NSX-T does not support dynamic address groups for North-South traffic.
6. Select the **Application** to allow. In this example, we create an **Application Group** that includes a static group of specific applications that are grouped together.
   1. Click **Add** and select **New Application Group**.
   2. Click **Add** to select the application to include in the group.
   3. Click **OK** to create the application group.
7. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles for antivirus, anti-spyware, and vulnerability protection, under Profiles.
8. Click **Commit**, select **Commit to Panorama**. Click **OK**.

STEP 3 | Apply the policies to the VM-Series firewalls on NSX-T.

1. Click **Commit** > **Push to Devices** > **Edit Selections**.
2. Select the device group and click **OK**.
3. Select **Force Template Values**. By default, Panorama does not override objects on the firewall with objects on Panorama that share a name. You must select Force Template Values to push policy to the managed firewalls.
4. Click **Yes** to confirm force template values.
5. Click **OK**.
6. Verify that the commit is successful.

STEP 4 | (Optional) Use template to push a base configuration for network and device configuration such as DNS server, NTP server, Syslog server, and login banner.

Refer to the Panorama Administrator's Guide for information on using templates.

Use vMotion to Move the VM-Series Firewall Between Hosts

To maintain traffic flow while using vMotion to move your VM-Series firewall between ESXi hosts with homogeneous CPU configurations in VMware NSX-T, you must use the PAN-OS CLI to pause the internal heartbeat monitoring of the VM-Series firewall during vMotion. You can

 ©2024 Palo Alto Networks, Inc.

specify the amount of time, in minutes, that heartbeat monitoring is paused. Heartbeat monitoring can be paused for up to 60 minutes. When the pause interval expires or you deliberately end the pause interval, heartbeat monitoring resumes.

vMotion of the VM-Series firewall is supported on vSphere 6.5, 6.7, and 7.0 if the ESXi hosts have homogeneous CPU configuration.

This procedure is not required when using vMotion to move the VM-Series firewall if you are running vSphere 7.0 or later.

STEP 1 | Log in the VM-Series firewall CLI.

STEP 2 | Set the heartbeat monitoring pause interval using the following command. The pause begins as soon as the command is executed. If vMotion is taking longer than expected, you can rerun this command to set a new, longer interval that starts when the command is executed again.

```
request system heartbeat-pause set interval <pause-time-in-minutes>
```

You can view the time remaining in pause interval using the following command.

```
request system heartbeat-pause show interval
```

STEP 3 | (Optional) If you complete vMotion before the pause interval has elapsed, you can end the pause by setting the interval to zero (0).

```
request system heartbeat-pause set interval 0
```

Extend Security Policy from NSX-V to NSX-T

If you are moving from an NSX-V deployment to an NSX-T deployment or combining an NSX-T deployment with an NSX-V deployment, you can extend your existing security policy from NSX-V to NSX-T without having to recreate the policy rules. This is achieved by leveraging your existing device groups and sharing them between the NSX-V and NSX-T service definitions. After migrating your policy to NSX-T, you can continue using the VM-Series for NSX-V or remove your NSX-V deployment.

STEP 1 | Install the Panorama Plugin for VMware NSX 3.2.0 or later. See the Panorama Plugin for VMware NSX 3.2.0 Release Notes before upgrading.

STEP 2 | Configure an NSX-T service definition for each NSX-V service definition in your deployment. Do not create new device groups; instead use your existing NSX-V device groups. Using the existing device groups allows you to apply the same security policy rules used on NSX-V to the VM-Series firewalls deployed on NSX-T. If you have policy that reference a particular zone, add the same template stack from your NSX-V service definition to your NSX-T service definition. Additionally, if your device group references a particular template, ensure that you select the template stack that includes the template referenced in the device group.

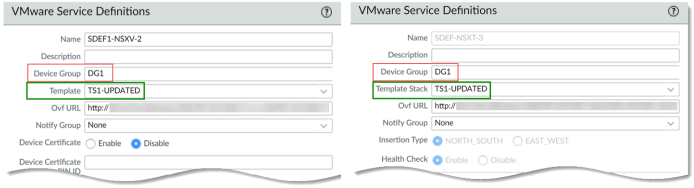

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure an NSX-T service manager and associate the NSX-T service definitions to the service manager.

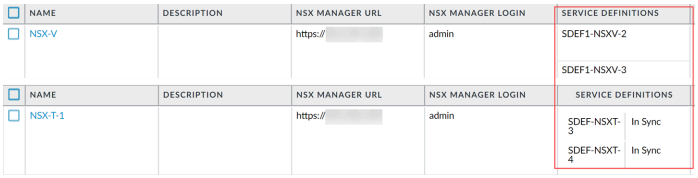

STEP 4 | Prepare your NSX-T environment and deploy the VM-Series firewall. You must create your security groups, service chains, and traffic redirection policy before launching the VM-Series firewall.

- Deploy the VM-Series Firewall on NSX-T (North-South)
- Deploy the VM-Series Using the Operations-Centric Workflow

STEP 5 | Add the NSX-T tags to you existing dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click on the name of an existing NSX-V dynamic address group.
3. Click **Add Match Criteria** to display the tags from NSX-V and NSX-T.
4. Add the NSX-T tag to the dynamic address groups. Be sure to use the **OR** operator between the tags.
5. When you have added all the necessary tags, click **OK**.
6. **Commit** your changes.

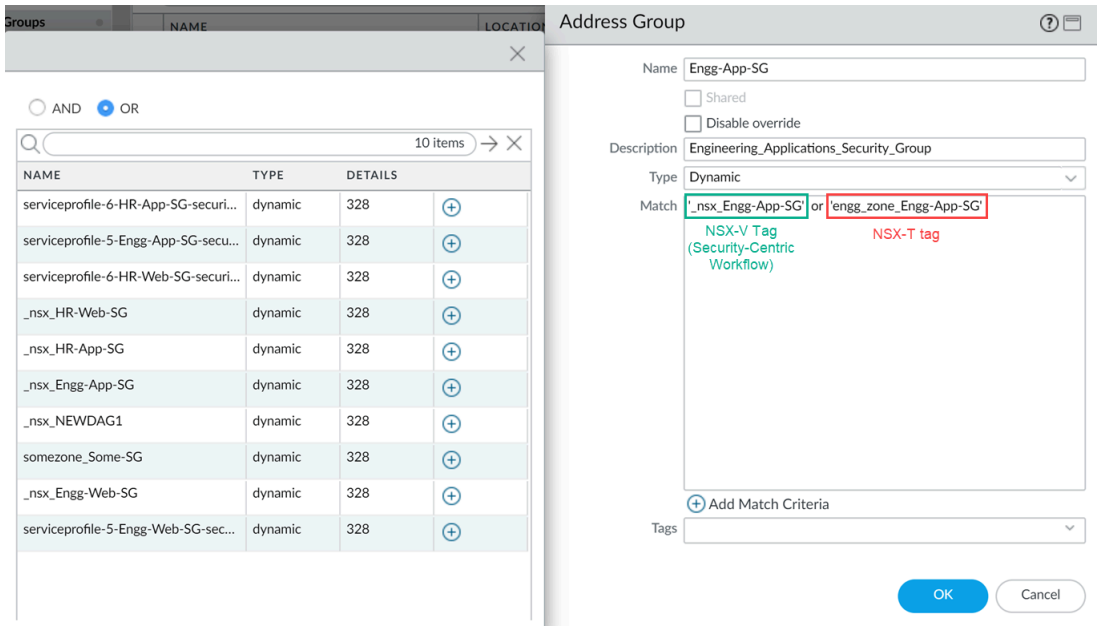

STEP 6 | After your VM workloads have successfully migrated from NSX-V to NSX-T, you remove the NSX-V tags from your dynamic address groups if you plan to discontinue use of NSX-V. All NSX-V tags and corresponding IP addresses are unregistered after all NSX-V related configuration is removed from the Panorama plugin for NSX and VM-Series firewall configuration is removed from NSX-V manager.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on NSX-T (East-West)

The VM-Series firewall on VMware NSX-T integrates the Palo Alto next-generation firewalls and Panorama with ESXi host servers to provide comprehensive visibility and safe application enablement of all East-West traffic in your NSX-T software-defined data center.

- Components of the VM-Series Firewall on NSX-T (East-West)
- VM-Series Firewall on NSX-T (East-West) Integration
- Supported Deployments of the VM-Series Firewall on VMware NSX-T (East-West)
- Deploy the VM-Series Using the Operations-Centric Workflow
- Deploy the VM-Series Using the Security-Centric Workflow
- Delete a Service Definition from Panorama
- Migrate from VM-Series on NSX-T Operation to Security Centric Deployment
- Extend Security Policy from NSX-V to NSX-T
- Use In-Place Migration to Move Your VM-Series from NSX-V to NSX-T
- Use Migration Coordinator to Move Your VM-Series from NSX-V to NSX-T

Components of the VM-Series Firewall on NSX-T (East-West)

The following tables show the components of this joint Palo Alto Networks and VMware NSX-T (East-West) solution.

VMware Components	
vCenter/ESXi	The vCenter server is the centralized management tool for the vSphere suite. ESXi is a hypervisor that enables compute virtualization. Refer to VMware's Compatibility Matrix for vCenter compatibility with your version of NSX-T.
NSX-T Manager	VMware NSX-T Data Center 2.5.0 and later must be installed and registered with the vCenter server. The NSX-T Manager is required to deploy the VM-Series firewall on the ESXi hosts within a ESXi cluster.

Palo Alto Networks Components	
PAN-OS	PAN-OS 9.1.x and later. The VM-Series base image, for example PA-VM-NST-9.1.0zip, is required for deploying the VM-Series firewall on NSX-T.

 ©2024 Palo Alto Networks, Inc.

Palo Alto Networks Components	
	The minimum system requirement for deploying the VM-Series firewall for NSX on the ESXi server depends on your VM-Series model. See VM-Series Models for the minimum hardware requirements for your VM-Series model.
Panorama Panorama must be running the same release version or later version that the firewalls that it will manage.	The VM-Series firewall on NSX-T requires Panorama 9.1.0 and later for firewalls running 9.1.0 Panorama is the centralized management tool for the Palo Alto Networks next-generation firewalls. In this solution, Panorama works with the NSX-T Manager to deploy, license, and centrally administer—configuration and policies—the VM-Series firewall for NSX-T. Panorama must be able to connect to the NSX-T Manager, the VM-Series firewalls and the Palo Alto Networks update server. See the 9.1 Panorama Administrator's Guide for information about deploying your Panorama appliance.
Panorama Plugin for VMware NSX	3.1.0 or later
VM-Series Plugin	1.0.8 or later
VM-Series Firewall Models	The VM-100, VM-300, VM-500, and VM-700 support NSX-T. *Before you deploy the VM-Series firewall on NSX-T, ensure that you have sufficient hardware resources to support the number of VM-Series firewalls in your chosen deployment model (service cluster or per host). This is critical when deploying large firewalls, such as the VM-700.*

VM-Series Firewall on NSX-T (East-West) Integration

NSX-T Manager, vCenter, Panorama, and the VM-Series firewall work together to meet the security challenges of your NSX-T Data Center.

 ©2024 Palo Alto Networks, Inc.

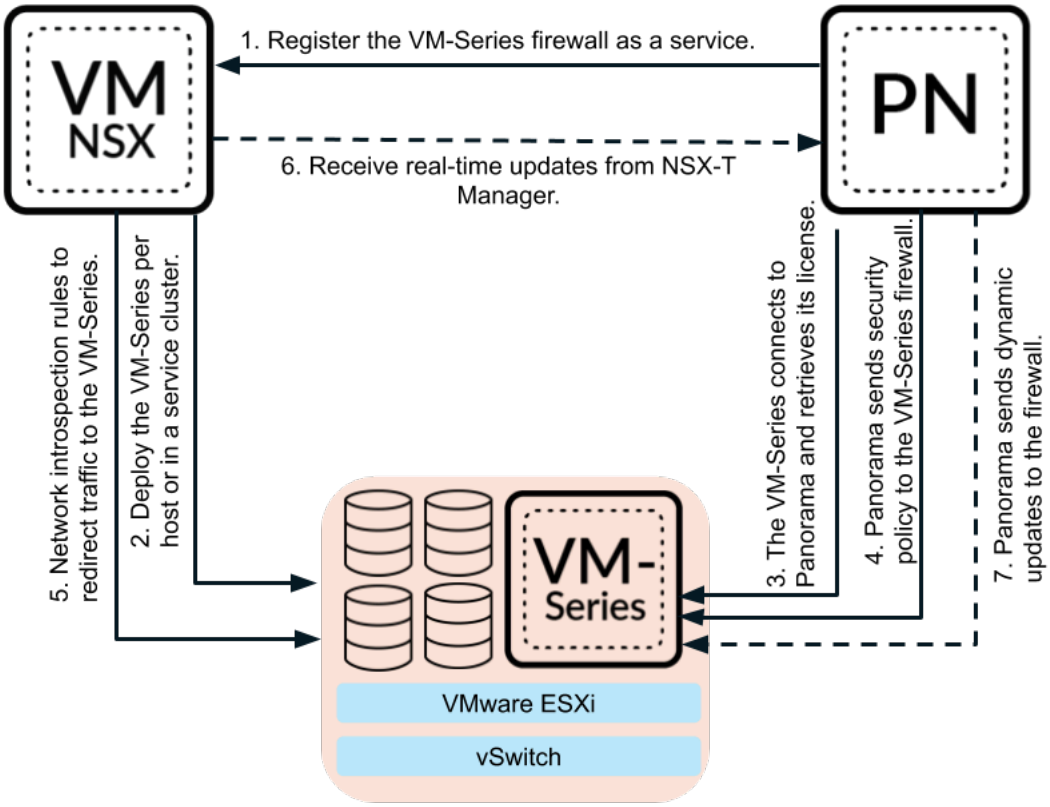

1. **Register the VM-Series firewall as a service**—Use Panorama to connect to your VMware NSX-T manager. Panorama communicates with NSX-T Manager using the NSX-T API and establishes bi-directional communication. On Panorama, you configure the Service Manager by entering the IP address, username, and password of NSX-T Manager to initiate communication.

   After establishing communication with NSX-T Manager, configure the service definition. The service definition includes the location of the VM-Series firewall base image, the authorization code needed to license the VM-Series firewall, and the device groups and template stack to which the firewall will belong.

   Additionally, NSX-T Manager uses this connection to send updates on the changes in the NSX-T environment with Panorama.
2. **Deploy the VM-Series firewall per host or in a service cluster**—NSX-T Manager uses the information pushed from Panorama in the service definition to deploy the VM-Series firewall. Choose a where the VM-Series firewall will be deployed (in a service cluster or on each ESXi host) and how NSX-T provides a management IP address to the VM-Series firewall (DHCP or static IP). When the firewall boots up, NSX-T manager's API connects the VM-Series firewall to the hypervisor so it that can receive traffic from the vSwitch.
3. **The VM-Series connects to Panorama**—The VM-Series firewall then connects to Panorama to obtain its license. Panorama gets the license from the Palo Alto Networks update server and

 ©2024 Palo Alto Networks, Inc.

sends it to the firewall. When the firewall gets its license, it reboots and comes back up with a serial number.

> *If Panorama does not have internet access, it cannot retrieve licenses and push them to the firewall, so you have to manually license each firewall individually. If the VM-Series firewall does not have internet access, you must manually add the serial numbers to Panorama to register them as managed devices, so Panorama can push template stacks, device groups, and other configuration information. For more information, see* Activate the License for the VM-Series Firewall for VMware NSX.

4. **Panorama sends security policy to the VM-Series firewall**—When the firewall reconnects to Panorama, it is added to device group and template stack defined in the service definition and Panorama pushes the appropriate security policy to that firewall. The firewall is now ready to secure traffic in your NSX-T data center.
5. **Create network introspection rules to redirect traffic to the VM-Series firewall**—On the NSX-T Manager, create a service chain and network introspection rules that redirect traffic in your NSX-T data center.
6. **Send real-time updates from NSX-T Manager**—The NSX-T Manager sends real-time updates about changes in the virtual environment to Panorama. These updates include changes in group membership and IP addresses of virtual machines in groups that send traffic to the VM-Series firewall.
7. **Panorama sends dynamic updates**—As Panorama receives updates from NSX-T Manager, it sends those updates from its managed VM-Series firewalls. Panorama places virtual machines into dynamic address groups based on criteria that you determine and pushes dynamic address group membership information to the firewalls. This allows firewalls to apply the correct security policy to traffic flowing to and from virtual machines in your NSX-T data center.

Supported Deployments of the VM-Series Firewall on VMware NSX-T (East-West)

You can deploy one or more instances of the VM-Series firewall as a partner service in your VMware NSX-T Data Center to secure East-West traffic and perform micro-segmentation. To configure the VM-Series firewall to perform mirco-segmentation, you can deploy the firewalls in a service cluster or per host.

- **Service Cluster**—In a clustered deployment, all the VM-Series firewalls are installed on a single cluster. Traffic between VMs and groups are redirected to the VM-Series cluster for policy inspection and enforcement before continuing to its destination. When you configure a

 ©2024 Palo Alto Networks, Inc.

clustered deployment, you can specify a particular host within the cluster or select **Any** and let NSX-T choose a host.

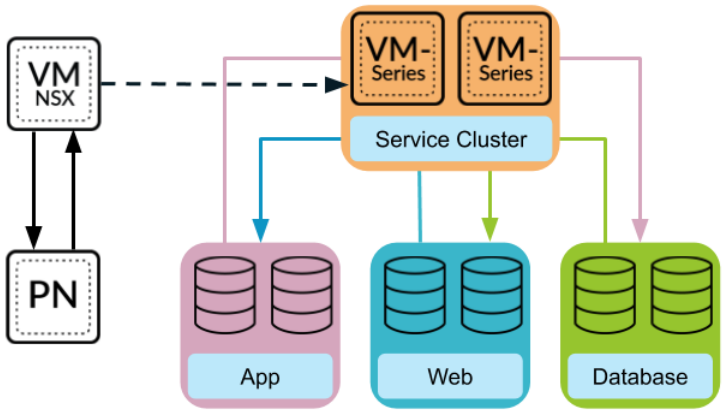

- **Host-Based**—In a per host deployment, an instance of the VM-Series firewall is installed on each host in the ESXi cluster. Traffic between guests on the same host is inspected by the local firewall, so it does not need to leave the host for inspection. Traffic leaving the host is inspected by the firewall before reaching the vSwitch.

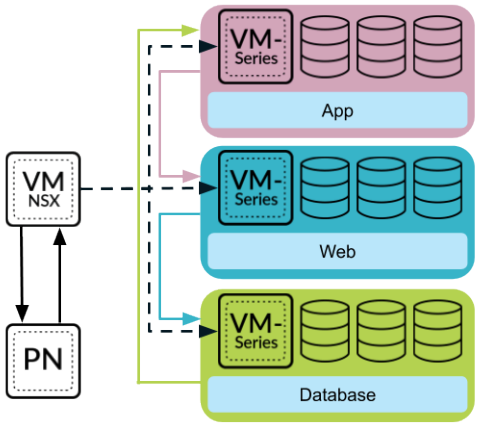

After deploying the firewall, you configure traffic redirection rules that send traffic to the VM-Series firewall. Security policy rules that you configure on Panorama are pushed to managed VM-Series firewalls and then applied to traffic passing through the firewall.

Deploy the VM-Series Using the Operations-Centric Workflow

Complete the following tasks to deploy the VM-Series firewall to secure East-West traffic in your NSX-T data center.

- Install the Panorama Plugin for VMware NSX
- Enable Communication Between NSX-T Manager and Panorama
- Create Template Stacks and Device Groups on Panorama
- Configure the Service Definition on Panorama
- Launch the VM-Series Firewall on NSX-T (East-West)
- Add a Service Chain

 ©2024 Palo Alto Networks, Inc.

- Direct Traffic to the VM-Series Firewall
- Apply Security Policies to the VM-Series Firewall on NSX-T (East-West)
- Use vMotion to Move the VM-Series Firewall Between Hosts

Install the Panorama Plugin for VMware NSX

Download and install the Panorama Plugin for VMware NSX. See the Compatibility Matrix before installing or upgrading your plugin.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panorama HA peers, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  request plugins dau plugin-name <plugin-name> unblock-device-push yes

  You can cancel this command by executing:

  request plugins dau plugin-name <plugin-name> unblock-device-push no

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Select **Panorama** > **Plugins**.

STEP 2 | Select **Check Now** to retrieve a list of available updates.

STEP 3 | Select **Download** in the Action column to download the plugin.

STEP 4 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

Enable Communication Between NSX-T Manager and Panorama

Complete the following procedure to enable communication between Panorama and NSX-T Manager. You can connect your Panorama to up to 16 NSX-T Managers. If you are connecting your Panorama to multiple NSX-T Managers, you must carefully plan your device group hierarchy and template stacks and consider how they interact with the other components needed for

 ©2024 Palo Alto Networks, Inc.

deployment. Service definitions reference device groups and template stacks and push that information to the firewalls in the related ESXi clusters.

STEP 1 | (Optional) Bypass proxy server settings, configured on Panorama under **Panorama** > **Setup** > **Services** > **Proxy Server**, for communication between Panorama and NSX-T Manager. This command allows Panorama to communicate directly with NSX-T Manager while maintaining proxied communication for other services.

1. Log in to the Panorama CLI.
2. Execute the following command to enable or disable proxy bypass.

   ```
   admin@Panorama> request plugins vmware_nsx proxy bypass {yes | no}
   ```

   Select **yes** to enable proxy bypass and **no** to disable proxy bypass. This is set to **no** by default.

STEP 2 | Log in to the Panorama web interface.

Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<*IP address*>).

STEP 3 | Set up access to the NSX-T Manager. Repeat this procedure for each NSX-T Manager to which you will connect Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers** and click **Add**.
2. Enter a descriptive **Name** for your NSX-T Manager.
3. (Optional) Add a **Description** for NSX-T Manager.
4. Enter the **NSX Manager URL**—NSX-T Manager cluster virtual IP address or FQDN—at which to access the NSX-T Manager.
5. Enter the **NSX Manager Login** credentials—username and password, so that Panorama can authenticate to the NSX-T Manager.
6. Click **OK**.

If you change your NSX-T Manager login password, ensure that you update the password on Panorama immediately. An incorrect password breaks the connection between Panorama and NSX-T Manager.

STEP 4 | Commit your changes to Panorama.

Select **Commit** and **Commit to Panorama**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Verify the connection status on Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers**.
2. Verify the message in the **Status** column.

   When the connection is successful, the status displays as **Registered**. This indicates that Panorama and the NSX-T Manager are in sync.

   The unsuccessful status messages are:

   - **No connection**: Unable to reach/establish a network connection to the NSX-T Manager.
   - **Invalid Credentials**: The access credentials (username and/or password) are incorrect.
   - **Out of sync**: The configuration settings defined on Panorama are different from what is defined on the NSX-T Manager.Click the link for details on the reasons for failure. For example, NSX-T Manager may have a service definition with the same name as defined on Panorama. To fix the error, use the service definition name listed in the error message to validate the service definition on the NSX-T Manager. Until the configuration on Panorama and the NSX-T Manager is synchronized, you cannot add a new service definition on Panorama.
   - **Connection Disabled**: The connection between Panorama and the NSX-T Manager was manually disabled.

Create Template Stacks and Device Groups on Panorama

To manage the VM-Series firewalls for NSX-T using Panorama, the firewalls must belong to a device group and a template that is a member of a template stack. Device groups allow you to assemble firewalls that need similar policies and objects as a logical unit; the configuration is defined using the **Objects** and **Policies** tabs on Panorama. Use template stacks to configure the settings that are required for the VM-Series firewalls to operate on the network and associate; the configuration is defined using the **Device** and **Network** tabs on Panorama. And each template stack with zones used in your NSX-T configuration on Panorama must be associated with a service definition; at a minimum, you must create a zone within the template stack so that the NSX-T Manager can redirect traffic to the VM-Series firewall.

Panorama can support deployments of both NSX-T North-South and NSX-T East-West at the same time. It is recommend that you configure separate device groups, template stacks, and service definitions for NSX-T North-South and NSX-T East-West.

STEP 1 | Add a device group or a device group hierarchy.

1. Select **Panorama** > **Device Groups**, and click **Add**. You can also create a device group hierarchy.
2. Enter a unique **Name** and a **Description** to identify the device group.
3. Click **OK**.
4. Click **Commit** and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Add a template.

1. Select **Panorama** > **Templates**, and click **Add**.
2. Enter a unique **Name** and a **Description** to identify the template.
3. Click **OK**.
4. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 3 | Create a template stack.

1. Select **Panorama** > **Templates**, and click **Add Stack**.
2. Enter a unique **Name** and a **Description** to identify the template.
3. Click **OK**.
4. Click **Commit**, and select **Commit to Panorama** to save the changes to the running configuration on Panorama.

STEP 4 | Create the zone(s) for each template.

Each zone is mapped to a service profile on NSX-T Manager. To qualify, a zone must be of the virtual wire type and a template associated with a service definition.

You can add up to 32 zones in each template.

1. Select **Network** > **Zones**.
2. Select the correct template in the **Template** drop-down.
3. Select **Add** and enter a zone **Name**.
4. Set the interface **Type** to **Virtual Wire**.
5. Click **OK**.
6. Verify that the zones are attached to the correct template.

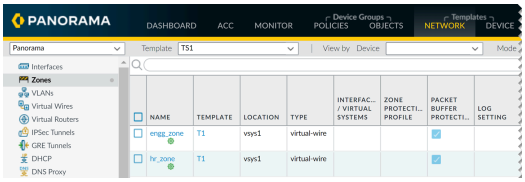

7. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

   Panorama creates a corresponding service profile on NSX-T Manager for each qualified zone upon commit.

Configure the Service Definition on Panorama

A service definition specifies the configuration for the VM-Series firewalls installed in your NSX-T data center environment. The service definition must include the device group, a template stack, and an OVF URL.

STEP 1 | (Optional) Configure a Notify Group

Create a notify group by specifying devices groups that should be notified of changes in the virtual environment. The firewalls included in the specified device groups receive a real-time update of security groups and IP addresses of guest VMs in them. The firewalls use this

 ©2024 Palo Alto Networks, Inc.

update to determine the most current list of members that constitute dynamic address groups referenced in policy

1. Select **Panorama** > **VMware** > **Notify Group** and click **Add**.
2. Give your Notify Group a descriptive **Name**.
3. Select the boxes of all devices groups that should be notified of changes to the virtual environment. If a device group does not have a check box available, it means that the device group is automatically included by virtue of the device group hierarchy.
4. Click **OK**.

STEP 2 | Add a new service definition.

You can create up to 32 service definitions on Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Definitions**.
2. Select **Add** to create a new service definition.
3. Enter a descriptive **Name** for your service definition.
4. (Optional) Add a **Description** that identifies the function or purpose for the VM-Series firewalls that will be deployed using this service definition.

STEP 3 | Assign a device group and a template stack to the service definition.

Make sure to Create Template Stacks and Device Groups on Panorama.

Because the firewalls deployed in this solution will be centrally administered from Panorama, you must specify the **Device Group** and the **Template Stack** that the firewalls belong to. All the firewalls that are deployed using this service definition belong to the specified template stack and device group.

1. Select the device group or device group hierarchy in the **Device Group** drop-down.
2. Select the template stack in the **Template** drop-down.

You cannot reuse a template stack or a device group assigned to one service definition in another service definition.

STEP 4 | Specify the location of the OVF file.

Download the zip file, unzip it to extract and save the .ovf, mf and .vmdk files to the same directory. The ovf and vmdk files are used to deploy each instance of the firewall.

Do not change the Panorama service definition OVF path after a successful NSX Service Deployment of VM-Series firewalls. Changing the OVF path, after a successful VM-Series firewall deployment, can result in a NSX Service Deployment failed state. You may resolve this failure in NSX-T Manager, however this may cause all VM-Series firewalls to redeploy.

It is recommended that you use an OVF path name that scales and allows you to change the base image without impacting your deployed firewalls. Instead of a path such as **https://acme.com/software/PA-VM-NST.9.1.0.ovf**, use something such as **https://acme.com/software/PanoSvcDef1-Cluster1.ovf**. Using a static path reference will eliminate any future

 ©2024 Palo Alto Networks, Inc.

need to change the OVF path. It is recommended to create a path for each Panorama service definition (vSphere cluster) in your deployment and change the PAN-OS base images references on the web server as needed.

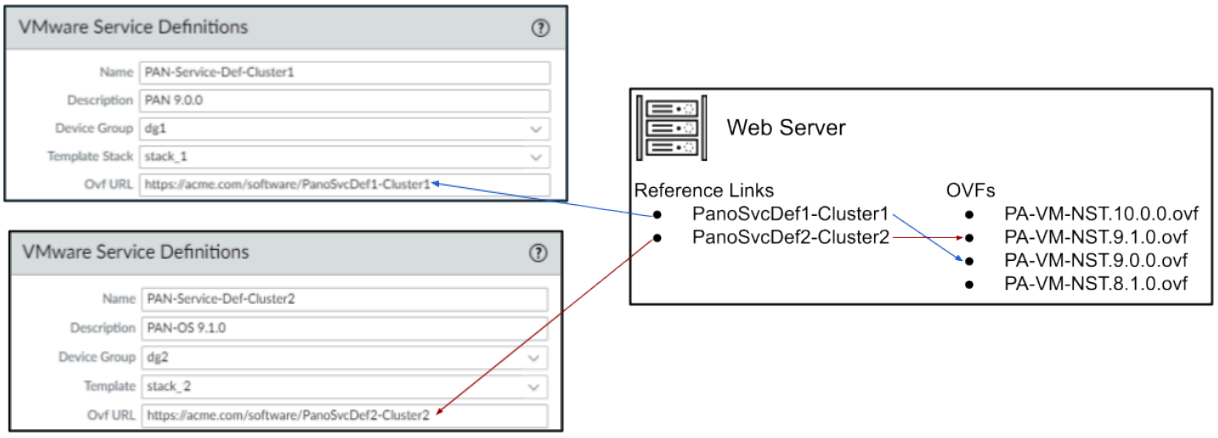

In **OVF URL**, add the location of the web server that hosts the ovf file. Both http and https are supported protocols.

You can use the same ovf version or different versions across service definitions. Using different ovf versions across service definitions allows you to vary the PAN-OS version on the VM-Series firewalls in different ESXi clusters.

STEP 5 | (Optional) Select a **Notify Group.**

STEP 6 | Select **East West** as the **Insertion Type** for your firewall.

STEP 7 | (Optional) **Enable Health Check**. Health check is disabled by default. Also called service health check, this NSX-T feature allows you to simulate high availability in the case of a service instance failing. When configured with the VM-Series firewall, if a VM-Series service instance fails, any traffic directed to that firewall is redirect to another firewall instance in the cluster (for service cluster deployments) or a firewall instance on another host (for host-based deployments).

You cannot disable or enable Health Check in a service definition after committing and deploying VM-Series firewalls in NSX-T. Attempting to commit a change in the Health Check configuration returns commit failure. To change this, you must delete and recreate your service definition and redeploy your VM-Series firewalls.

STEP 8 | To automatically retrieve a device certificate when the VM-Series firewall is deployed by NSX Manager, configure the device certificate.

Enable this option to apply a device certificate to newly deployed VM-Series firewalls. Only use this option when deploying the firewall using a base image OVF that supports device certificates. Panorama pushes the device certificate information to NSX Manager as part of the

 ©2024 Palo Alto Networks, Inc.

service definition. When a new firewall is deployed in NSX, the device certificate is installed on the firewall at bootup.

For list of OVFs that support device certificates for the VM-Series firewall on VMware NSX, see the Palo Alto Networks Compatibility Matrix.

If your OVF does not support a device certificate, disable this option.

1. If you have not done so already, log in to the Customer Support Portal and generate a Registration PIN and PIN ID.
2. Under **Device Certificate**, click **Enable**.
3. Copy the PIN ID and enter it into the **Device Certificate PIN ID** field.
4. Reenter the PIN ID into the **Confirm Device Certificate PIN ID** field.
5. Copy the PIN Value and enter it into the **Device Certificate PIN Value** field.
6. Reenter the PIN Value into the **Confirm Device Certificate PIN Value** field.

STEP 9 | Click **OK** to save the service definition.

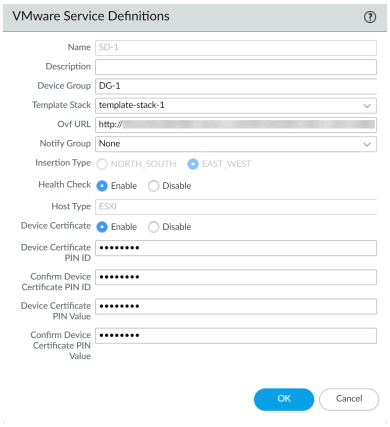

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Attach the service definition to the service manager.

You cannot use a service definition in more than one service manager.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Manager** and click the link of the service manager name.
2. Under Service Definitions, click **Add** and select your service definition from the drop-down.
3. Click **OK**.

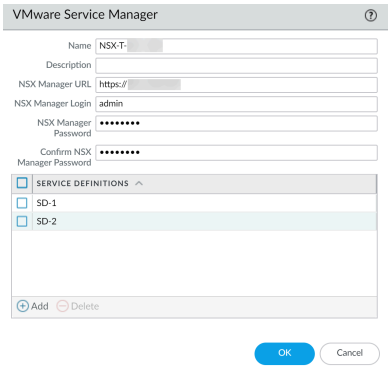

STEP 11 | Add the authorization code to license the firewalls.

1. Select **Panorama** > **Device Groups** and choose the device group you associated with the service definition you just created.
2. Under **Dynamically Added Device Properties**, add the authorization code you received with your order fulfillment email and, optionally, select None from the **SW Version** drop-down.

   When a new firewall is deployed on NSX-T it is automatically added to the device group, licensed using the authorization code you provided, and upgraded to the PAN-OS version you specified.

   On the support portal, you can view the total number of firewalls that you are authorized to deploy and the ratio of the number of licenses that have been used to the total number of licenses enabled by your authorization code.

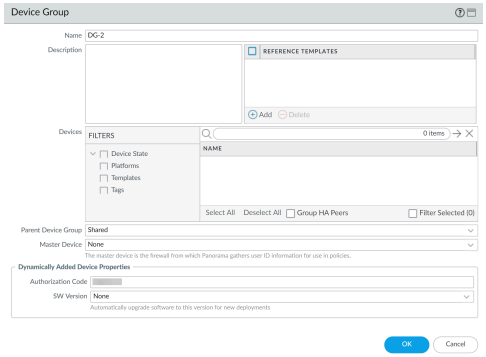

STEP 12 | **Commit to Panorama**.

 ©2024 Palo Alto Networks, Inc.

STEP 13 | On the NSX-T Manager, verify that the service definition is available.

Select **System** > **Service Deployments** > **Catalog**. The service definition is listed as a Service Instance on the NSX-T Manager.

Launch the VM-Series Firewall on NSX-T (East-West)

Complete the following procedure to deploy the VM-Series firewall as a service in your NSX-T environment. The **Deployment Specification** and **Deployment Template** fields are automatically populated with information pushed from Panorama as part of the service definition.

Do not edit any settings under Deployment Attributes. These values are imported from Panorama and changing them causes the deployment to fail.

STEP 1 | Log in to the NSX-T Manager.

STEP 2 | Select **System** > **Service Deployments** > **Deployment**.

STEP 3 | Select your service definition from the **Partner Service** drop-down.

STEP 4 | Click **Deploy Service**.

STEP 5 | Enter a descriptive **Name** for your service deployment.

STEP 6 | Select the **Compute Manager** (vCenter).

STEP 7 | Select a **Deployment Type**—**Clustered** or **Host Based**.

STEP 8 | If you selected **Clustered** as the **Deployment Type**, enter the **Clustered Deployment Count** to specify the number of VM-Series firewall instances to deploy on the cluster.

STEP 9 | Select a **Host** if you are launching the VM-Series in a clustered deployment. Select a particular host from the **Host** drop-down or **Any** to allow NSX-T Manager to choose the host. This option is grayed out in **Per Host** deployments.

STEP 10 | Select a **Data Store** as the repository for the VM-Series firewall. In a clustered deployment, select a shared data store if you choose **Any** for the host or select a local data store if you specified a particular host.

STEP 11 | Configure the **Networks** settings.

1. In the Networks column, click **Set**.
2. Select the **Network** for **eth0 - Management Nic**.
3. Select the **Network Type**—DHCP or Static IP Pool. If you choose Static IP Pool, select an **IP Pool**.
4. Check **eth1 - Data-1 Nic**.
5. Click **Save**.

 ©2024 Palo Alto Networks, Inc.

STEP 12 | Select or configure a **Service Segment**. To configure a service segment, complete the following procedure.

1. Click **Action** in the **Service Segments** column.

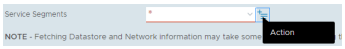

2. Click **Add Service Segment**.
3. Enter a descriptive **Name**.
4. Select a **Transport Zone (Overlay)**.
5. Click **Save** and **Close**.

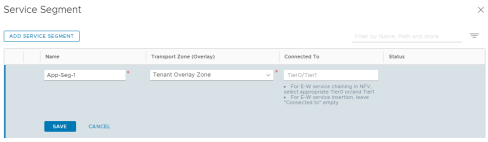

STEP 13 | Select the **Cluster** where the service will be deployed. You must select a cluster with **NSX Configuration**.

STEP 14 | Click **Save**.

STEP 15 | Verify that your firewalls deployed successfully.

1. Select **System** > **Service Deployments** > **Service Instances**.
2. Confirm that your firewalls are listed and the **Deployment Status** shows **Up**.

STEP 16 | Verify that your firewalls connected to Panorama.

1. Log in to Panorama.
2. Select **Panorama** > **Managed Devices** > **Summary**.
3. Confirm that your firewalls are listed under the correct device group and the **Device State** shows **Connected**.

   The Device Name for the VM-Series firewall is displayed on Panorama as **PA-VM:<nsx.clusterid>** for NSX-T (N-S) deployment and as **PA-VM:<nsx.servicevmid>** for NSX-T (E-W) deployment.

STEP 17 | Set a secure password for the admin account on your VM-Series firewalls.

Each VM-Series firewall uses a default username and password (admin/admin), which is used for initial login. Upon logging in for the first time, you are prompted to set a new, more secure

 ©2024 Palo Alto Networks, Inc.

password. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

You can update the password on each firewall individually or all at once through Panorama.

- **Panorama**—on Panorama, you can change the default password for all firewalls in a template or delete the admin user and create a new username and password.
  1. Log in to Panorama
  2. Select **Device** > **Administrators** and select the **admin** user.
  3. **Delete** the user or click the user and enter a new password.
  4. If you changed the password, click **OK**.
  5. Select **Commit** > **Push to Devices** > **Edit Selections** > **Force Template Values**.
  6. Click **OK**.
- **Firewall**—this procedure must be repeated on each VM-Series firewall.
  1. Log in to the VM-Series firewall using the default username and password.
  2. Follow the prompts to reset the password.

Add a Service Chain

A service chain is a grouping of services set in logical sequence. When traffic is redirected to the service chain, it moves through each service in the order you configure.

STEP 1 | Select **Security** > **Network Introspection Settings** > **Service Chains** > **Add Chain**.

STEP 2 | Enter a descriptive **Name** and **Description** (optional) for your service chain.

STEP 3 | Select the **Service Segment** that you applied when you deployed the VM-Series firewall.

STEP 4 | Set the forward path. The service chain is a logical sequence of service profiles, so traffic moves through the services in the order you specify as the forward path.

1. Select **Set Forward Path** > **Add Profile in Sequence**.
2. Select a service profile. The service column is populated automatically based on the service profile you select.
3. Click **Add**.
4. (Optional) If you have other partner service profiles in your NSX-T environment, click **Add Profile in Sequence** to add them to this service chain.

   You can select only one service profile per service definition.

5. Click **Save** when you have finished adding service profiles.

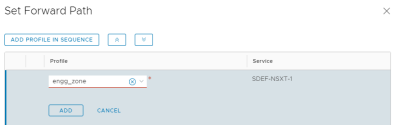

STEP 5 | In the Reverse Path column, check **Inverse ForwardPath** for return traffic to move through the service chain in reverse order.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | (Optional) If other partner service profiles are selected, set a reverse path.

You must select the same VM-Series service profile set in the Forward Path.

1. Select **Set Reverse Path** > **Add Profile in Sequence**.
2. Select a service profile. The service column is populated automatically based on the service profile you select.
3. Click **Add**.
4. (Optional) If you have other, service profiles in your NSX-T environment, click **Add Profile in Sequence** to add them to this service chain.
5. Click **Save** when you have finished adding service profiles.

STEP 7 | Set the **Failure Policy**—**Allow** or **Block**. This defines the action NSX-T takes if a service profile fails.

STEP 8 | Click **Save**.

Direct Traffic to the VM-Series Firewall

Configure policy rules to direct traffic virtual machines or groups of virtual machines to the VM-Series firewall.

STEP 1 | Select **Security** > **Network Introspection (E-W)** > **Rules** > **Add Policy**.

STEP 2 | Click **New Policy** to give your policy a descriptive name.

STEP 3 | Select your service chain from the **Redirect To** drop-down.

STEP 4 | Select the policy and click **Add Rule**.

STEP 5 | Click **New Rule** to give your rule a descriptive name.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Select a source.

1. Click the pencil icon in the source column to choose a source group of virtual machines.

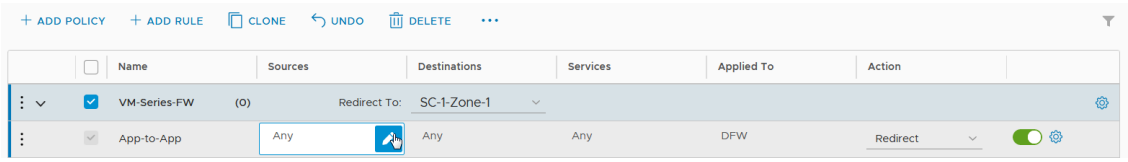

2. Check the source group or groups.
3. Click **Apply**.

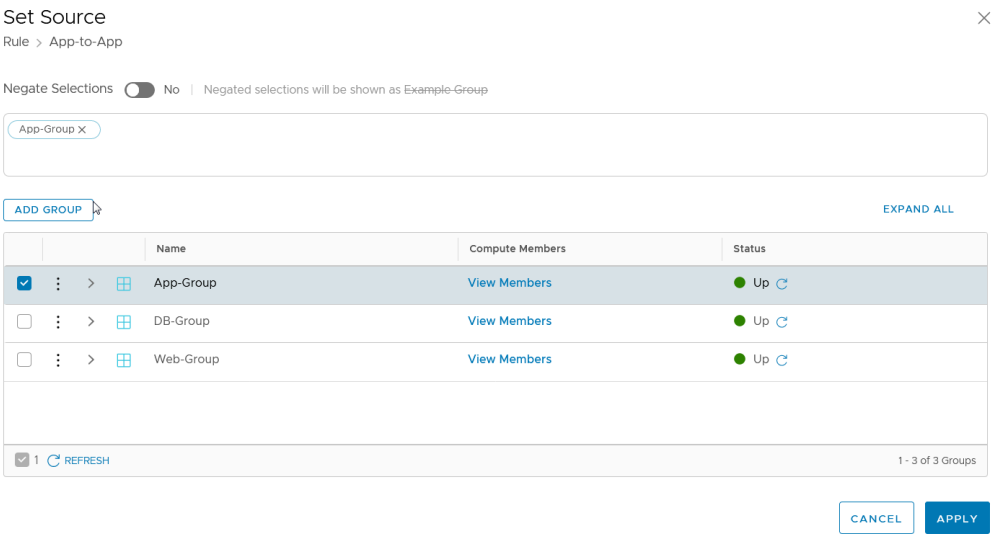

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Select a **Destination**.

1. Click the pencil icon in the destination column to choose a source group of virtual machines.

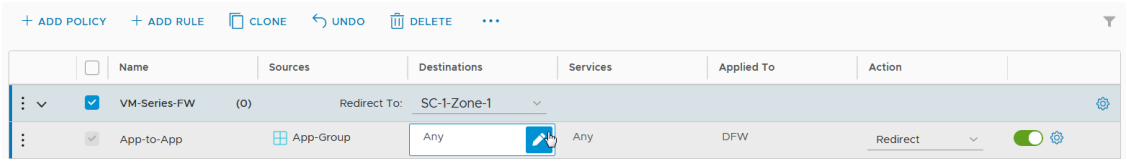

2. Check the destination group or groups.
3. Click **Apply**.

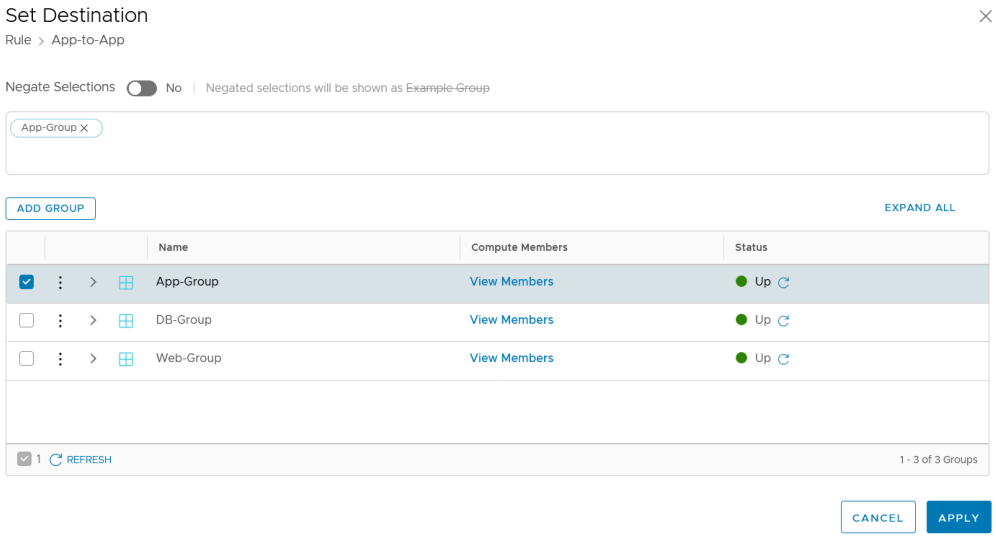

STEP 8 | (Optional) Select **Services** to which the rule will be applied.

STEP 9 | Choose one of the following in the **Applied To** field:

- Select **DFW** to apply the rule to all virtual NICs attached to the logical switch.
- Select **Groups** to apply the rule to virtual NICs of members virtual machines in the specified group or groups.

STEP 10 | Select the **Action**—**Redirect** or **Do Not Redirect**.

STEP 11 | Click **Publish**.

STEP 12 | Repeat this process to create additional policy or rules.

Apply Security Policies to the VM-Series Firewall on NSX-T (East-West)

Now that you have created the redirection rules on the NSX-T Manager, you can now use Panorama for centrally administering policies on the VM-Series firewalls.

To manage centralized policy, attach the dynamic address group as a source or destination address in security policy and push it to the firewalls; the firewalls can dynamically retrieve the IP addresses of the virtual machines that are included in each security group to enforce compliance for traffic that originates from or is destined to the virtual machines in the specified group.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Log in to Panorama.

STEP 2 | Create dynamic-address groups.

1. Select **Objects** > **Address Groups**.
2. Select the Device Group you created for managing your VM-Series on NSX-T firewall from the **Device Group** drop-down.
3. Click **Add** and enter a **Name** and **Description** for the dynamic address group.
4. Select **Type** as **Dynamic**.
5. Add Match Criteria to your dynamic address group.

   Some browser extensions may block API calls between Panorama and NSX-T which prevents Panorama from receiving match criteria. If Panorama displays no match criteria and you are using browser extensions, disable the extensions and Synchronize Dynamic Objects to populate the tags available to Panorama.
6. Click **Add Match Criteria**.
7. Select the **And** or **Or** operator and click the plus (+) icon next to the security group name to add it to the dynamic address group.

   The security groups that display in the match criteria dialog are derived from the groups you defined on the NSX-T Manager. Only the groups that are referenced in the security policies and from which traffic is redirected to the VM-Series firewall are available here.

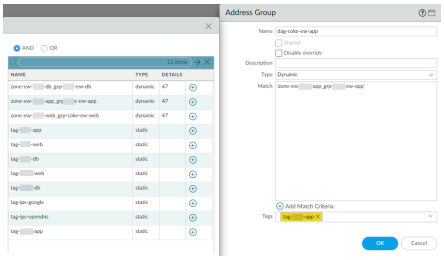

8. Click **OK**.
9. Repeat these steps to create the appropriate number of dynamic address groups required for your deployment.
10. **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create security policy rules.

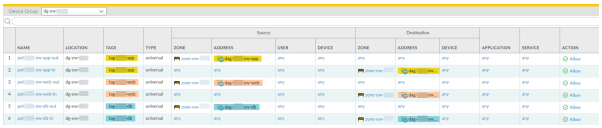

1. Select **Policies** > **Security** > **Prerules**.
2. Select the **Device Group** that you created for managing the VM-Series firewalls on NSX-T in Register the VM-Series Firewall as a Service on the NSX-V Manager.
3. Click **Add** and enter a **Name** and a **Description** for the rule. In this example, the security rule allows all traffic between the WebFrontEnd servers and the Application servers.
4. Select the **Source Zone** and **Destination Zone**. The zone name must be the same in both columns.
5. For the **Source Address** and **Destination Address**, select or type in an address, address group or region. In this example, we select an address group, the Dynamic address group you created previously.
6. Select the **Application** to allow. In this example, we create an **Application Group** that includes a static group of specific applications that are grouped together.
   1. Click **Add** and select **New Application Group**.
   2. Click **Add** to select the application to include in the group.
   3. Click **OK** to create the application group.
7. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles for antivirus, anti-spyware, and vulnerability protection, under Profiles.
8. Repeats the steps above to create the pertinent policy rules.
9. Click **Commit**, select Commit Type as **Panorama**. Click **OK**.

STEP 4 | Apply the policies to the VM-Series firewalls for NSX-T.

1. Click **Commit**, and select Commit Type **Device Groups**.
2. Select the device group, NSX-T Device Group in this example and click **OK**.
3. Verify that the commit is successful.

STEP 5 | Validate that the members of the dynamic address group are populated on the VM-Series firewall.

1. From Panorama, switch device context to launch the web interface of a firewall to which you pushed policies.
2. On the VM-Series firewall, select **Policies** > **Security**, and select a rule.
3. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.

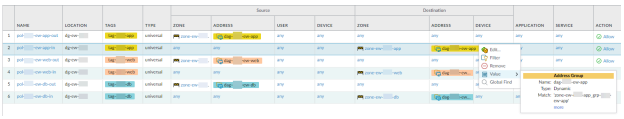

4. Click the **more** link and verify that the list of registered IP addresses is displayed.

   Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | (Optional) Use template to push a base configuration for network and device configuration such as DNS server, NTP server, Syslog server, and login banner.

Refer to the Panorama Administrator's Guide for information on using templates.

STEP 7 | Create a Zone Protection profile and attach it to a zone.

A zone protection profile provides flood protection and has the ability to protect against port scanning, port sweeps and packet-based attacks. It allows you to secure intra-tier and inter-tier traffic between virtual machines within your data center and traffic from the Internet that is destined to the virtual machines (workloads) in your data center.

1. Select your **Template**.
2. Select **Network** > **Network Profiles** > **Zone Protection** to add and configure a new profile.
3. Select **Network** > **Zones**, click the default-zone listed and select the profile in the **Zone Protection Profile** drop down.

STEP 8 | Create a DoS Protection profile and attach it to DoS Protection policy rule.

1. Select your **Device Group**.
2. Select **Objects** > **Security Profiles** > **DoS Protection** to add and configure a new profile.
   - A classified profile allows the creation of a threshold that applies to a single source IP. For example, you can configure a max session rate for an IP address that matched the policy, and then block that single IP address once the threshold is triggered.
   - An aggregate profile allows the creation of a max session rate for all packets matching the policy. The threshold applies to new session rate for all IP addresses combined. Once the threshold is triggered it affects all traffic that matches the policy.
3. Create a new DoS Protection policy rule in **Policy** > **DoS Protection,** and attach the new profile to it.

Use vMotion to Move the VM-Series Firewall Between Hosts

To maintain traffic flow while using vMotion to move your VM-Series firewall between ESXi hosts with homogeneous CPU configurations in VMware NSX-T, you must use the PAN-OS CLI to pause the internal heartbeat monitoring of the VM-Series firewall during vMotion. You can specify the amount of time, in minutes, that heartbeat monitoring is paused. Heartbeat monitoring can be paused for up to 60 minutes. When the pause interval expires or you deliberately end the pause interval, heartbeat monitoring resumes.

vMotion of the VM-Series firewall is supported on vSphere 6.5, 6.7, and 7.0 if the ESXi hosts have homogeneous CPU configuration.

This procedure is not required when using vMotion to move the VM-Series firewall if you are running vSphere 7.0 or later.

STEP 1 | Log in the VM-Series firewall CLI.

STEP 2 | Set the heartbeat monitoring pause interval using the following command. The pause begins as soon as the command is executed. If vMotion is taking longer than expected, you can

 ©2024 Palo Alto Networks, Inc.

rerun this command to set a new, longer interval that starts when the command is executed again.

```
request system heartbeat-pause set interval <pause-time-in-minutes>
```

You can view the time remaining in pause interval using the following command.

```
request system heartbeat-pause show interval
```

STEP 3 | (Optional) If you complete vMotion before the pause interval has elapsed, you can end the pause by setting the interval to zero (0).

```
request system heartbeat-pause set interval 0
```

Deploy the VM-Series Using the Security-Centric Workflow

You can use the security-centric workflow to control and manage your VM-Series firewall for NSX-T from Panorama. You do not need to access NSX-T Manager to create service chains and steering rules; however, the service deployment must still be created on NSX-T Manager.

- Install the Panorama Plugin for VMware NSX
- Enable Communication Between NSX-T Manager and Panorama
- Create Template Stacks and Device Groups on Panorama
- Configure the Service Definition on Panorama
- Launch the VM-Series Firewall on NSX-T (East-West)
- Create Dynamic Address Groups
- Create Security Policies
- Create Dynamic Address Group Membership Criteria
- Generate Steering Policy
- Generate Steering Rules

Install the Panorama Plugin for VMware NSX

Download and install the Panorama Plugin for VMware NSX. See the Compatibility Matrix before installing or upgrading your plugin.

The security-centric deployment workflow requires Panorama plugin for VMware NSX 4.0.0. Additionally, you must upgrade to plugin 4.0.0 from Panorama plugin for VMware NSX 3.2.x

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panorama HA peers, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

 ©2024 Palo Alto Networks, Inc.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  request plugins dau plugin-name <plugin-name> unblock-device-push yes

  You can cancel this command by executing:

  request plugins dau plugin-name <plugin-name> unblock-device-push no

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Select **Panorama** > **Plugins**.

STEP 2 | Select **Check Now** to retrieve a list of available updates.

STEP 3 | Select **Download** in the Action column to download the plugin.

STEP 4 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

Enable Communication Between NSX-T Manager and Panorama

Complete the following procedure to enable communication between Panorama and NSX-T Manager. You can connect your Panorama to up to 16 NSX-T Managers. If you are connecting your Panorama to multiple NSX-T Managers, you must carefully plan your device group hierarchy and template stacks and consider how they interact with the other components needed for deployment. Service definitions reference device groups and template stacks and push that information to the firewalls in the related ESXi clusters.

STEP 1 | (Optional) Bypass proxy server settings, configured on Panorama under **Panorama** > **Setup** > **Services** > **Proxy Server**, for communication between Panorama and NSX-T Manager.

This command allows Panorama to communicate directly with NSX-T Manager while maintaining proxied communication for other services.

1. Log in to the Panorama CLI.
2. Execute the following command to enable or disable proxy bypass.

   admin@Panorama> request plugins vmware_nsx proxy bypass {yes | no}

   Select **yes** to enable proxy bypass and **no** to disable proxy bypass. This is set to **no** by default.

STEP 2 | Log in to the Panorama web interface.

Using a secure connection (https) from a web browser, log in using the IP address and password you assigned during initial configuration (https://<*IP address*>).

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Set up access to the NSX-T Manager.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers** and click **Add**.
2. Enter a descriptive **Name** for your NSX-T Manager.
3. (Optional) Add a **Description** for NSX-T Manager.
4. Enter the **NSX Manager URL**—NSX-T Manager cluster virtual IP address or FQDN—at which to access the NSX-T Manager.
5. Enter the **NSX Manager Login** credentials—username and password, so that Panorama can authenticate to the NSX-T Manager.
6. Click **OK**.
7. Repeat this procedure for each NSX-T Manager to which you will connect Panorama.

If you change your NSX-T Manager login password, ensure that you update the password on Panorama immediately. An incorrect password breaks the connection between Panorama and NSX-T Manager.

STEP 4 | Commit your changes to Panorama.

Select **Commit** and **Commit to Panorama**.

STEP 5 | Verify the connection status on Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers**.
2. Verify the message in the **Status** column.

   When the connection is successful, the status displays as **Registered**. This indicates that Panorama and the NSX-T Manager are in sync.

   The unsuccessful status messages are:

   - **No connection**: Unable to reach/establish a network connection to the NSX-T Manager.
   - **Invalid Credentials**: The access credentials (username and/or password) are incorrect.
   - **Out of sync**: The configuration settings defined on Panorama are different from what is defined on the NSX-T Manager.Click the link for details on the reasons for failure. For example, NSX-T Manager may have a service definition with the same name as defined on Panorama. To fix the error, use the service definition name listed in the error message to validate the service definition on the NSX-T Manager. Until the configuration on Panorama and the NSX-T Manager is synchronized, you cannot add a new service definition on Panorama.
   - **Connection Disabled**: The connection between Panorama and the NSX-T Manager was manually disabled.

Create Template Stacks and Device Groups on Panorama

To manage the VM-Series firewalls for NSX-T using Panorama, the firewalls must belong to a device group and a template that is a member of a template stack. Device groups allow you to assemble firewalls that need similar policies and objects as a logical unit; the configuration is defined using the **Objects** and **Policies** tabs on Panorama. Use template stacks to configure the settings that are required for the VM-Series firewalls to operate on the network and associate; the configuration is defined using the **Device** and **Network** tabs on Panorama. And each template

 ©2024 Palo Alto Networks, Inc.

stack with zones used in your NSX-T configuration on Panorama must be associated with a service definition that you will create later; at a minimum, you must create a zone within the template stack so that the NSX-T Manager can redirect traffic to the VM-Series firewall. Later, you will associate a device group and template to your NSX-T deployment create a service definition.

Panorama can support deployments of both NSX-T North-South and NSX-T East-West at the same time. You must configure separate device groups, template stacks, and service definitions for NSX-T North-South and NSX-T East-West.

STEP 1 | Add a device group or a device group hierarchy.

1. Select **Panorama** > **Device Groups**, and click **Add**. You can also create a device group hierarchy.
2. Enter a unique **Name** and a **Description** to identify the device group.
3. Click **OK**.
4. Click **Commit** and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 2 | Add a template.

1. Select **Panorama** > **Templates**, and click **Add**.
2. Enter a unique **Name** and a **Description** to identify the template.
3. Click **OK**.
4. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 3 | Create a template stack and add your newly created template.

1. Select **Panorama** > **Templates**, and click **Add Stack**.
2. Enter a unique **Name** and a **Description** to identify the template stack.
3. Under **Templates**, click **Add** and select the template you created in step 2 from the drop-down.
4. Click **OK**.
5. Click **Commit**, and select **Commit to Panorama** to save the changes to the running configuration on Panorama.

STEP 4 | Create the zone(s) for each template.

The Panorama plugin for VMware NSX maps each zone to a service profile on NSX-T Manager. To qualify, a zone must be of the virtual wire type and be part of a template you will associate with a service definition; see Configure the Service Definition on Panorama for more

 ©2024 Palo Alto Networks, Inc.

information. In most uses cases, a single zone is sufficient. However, you must create multiple zones for multi-tenancy

You can add up to 32 zones in each template.

1. Select **Network** > **Zones**.
2. Select the correct template in the **Template** drop-down.
3. Select **Add** and enter a zone **Name**.
4. Set the interface **Type** to **Virtual Wire**.
5. Click **OK**.
6. Verify that the zones are attached to the correct template.

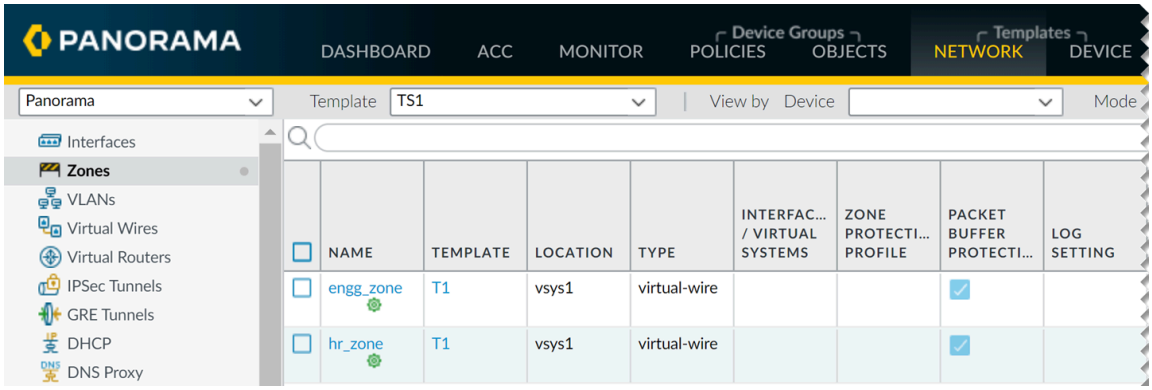

7. Click **Commit**, and select **Panorama** as the **Commit Type** to save the changes to the running configuration on Panorama.

STEP 5 | Update the DNS and NTP server information of your template stack. You must complete this step if you are using device certificates in your deployment. This is required to ensure the firewalls deployed in your NSX-T environment have the correct DNS information needed to reach the device certificate server.

1. Verify that you specified the correct template stack from the **Template** drop-down.
2. Select **Device** > **Setup** > **Services** and click the **Edit** icon.
3. On the Services tab, enter the IP address of the **Primary DNS Server** and **Secondary DNS Server**.
4. On the NTP tab, enter the IP address of the **NTP Server**.
5. Click **OK**.
6. **Commit** your changes to Panorama.

Configure the Service Definition on Panorama

A service definition allows you to register the VM-Series firewall as a partner security service on the NSX-T Manager. The service definition must include the device group, a template stack, and an OVF URL.

STEP 1 | (Optional) Configure a Notify Group

Create a notify group by specifying devices groups that should be notified of changes in the virtual environment. The firewalls included in the specified device groups receive a real-time update of security groups and IP addresses of guest VMs in them. The firewalls use this

 ©2024 Palo Alto Networks, Inc.

update to determine the most current list of members that constitute dynamic address groups referenced in policy.

1. Select **Panorama** > **VMware** > **Notify Group** and click **Add**.
2. Give your Notify Group a descriptive **Name**.
3. Select the boxes of all devices groups that should be notified of changes to the virtual environment. If a device group does not have a check box available, it means that the device group is automatically included by virtue of the device group hierarchy.
4. Click **OK**.

STEP 2 | Add a new service definition.

You can create up to 32 service definitions on Panorama.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Definitions**.
2. Select **Add** to create a new service definition.
3. Enter a descriptive **Name** for your service definition.
4. (Optional) Add a **Description** that identifies the function or purpose for the VM-Series firewalls that will be deployed using this service definition.

STEP 3 | Assign a device group and a template stack to the service definition.

Make sure to Create Template Stacks and Device Groups on Panorama.

Because the firewalls deployed in this solution will be centrally administered from Panorama, you must specify the **Device Group** and the **Template Stack** that the firewalls belong to. All the firewalls that are deployed using this service definition belong to the specified template stack and device group.

1. Select the device group or device group hierarchy in the **Device Group** drop-down.
2. Select the template stack in the **Template** drop-down.

You cannot reuse a template stack or a device group assigned to one service definition in another service definition.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Specify the location of the OVF file.

Download the zip file, unzip it to extract and save the .ovf, mf and .vmdk files to the same directory. The ovf and vmdk files are used to deploy each instance of the firewall.

Do not change the Panorama service definition OVF path after a successful NSX Service Deployment of VM-Series firewalls. Changing the OVF path, after a successful VM-Series firewall deployment, can result in a NSX Service Deployment failed state. You may resolve this failure in NSX-T Manager, however this may cause all VM-Series firewalls to redeploy.

In **OVF URL**, add the location of the web server that hosts the ovf file. Both http and https are supported protocols.

You can use the same ovf version or different versions across service definitions. Using different ovf versions across service definitions allows you to vary the PAN-OS version on the VM-Series firewalls in different ESXi clusters.

STEP 5 | (Optional) Select a **Notify Group.**

STEP 6 | Select **East West** as the **Insertion Type** for your firewall.

STEP 7 | (Optional) **Enable Health Check**.

Health check is enabled by default. Also called service health check, this NSX-T feature allows you to simulate high availability in the case of a service instance failing. When configured with the VM-Series firewall, if a VM-Series service instance fails, any traffic directed to that firewall is redirect to another firewall instance in the cluster (for service cluster deployments) or a firewall instance on another host (for host-based deployments).

You cannot disable or enable Health Check in a service definition after committing and deploying VM-Series firewalls in NSX-T. Attempting to commit a change in the Health Check configuration returns commit failure. To change this, you must delete and recreate your service definition and redeploy your VM-Series firewalls.

STEP 8 | To automatically retrieve a device certificate when the VM-Series firewall is deployed by NSX Manager, configure the device certificate.

Enable this option to apply a device certificate to newly deployed VM-Series firewalls. Only use this option when deploying the firewall using a base image OVF that supports device certificates. Panorama pushes the device certificate information to NSX Manager as part of the

 ©2024 Palo Alto Networks, Inc.

service definition. When a new firewall is deployed in NSX, the device certificate is installed on the firewall at bootup.

For list of OVFs that support device certificates for the VM-Series firewall on VMware NSX, see the Palo Alto Networks Compatibility Matrix.

If your OVF does support a device certificate, you must Enable device certificates regardless of whether or not you are using a device certificate. If your OVF does not support a device certificate, disable this option.

1. If you have not done so already, log in to the Customer Support Portal and generate a Registration PIN and PIN ID.
2. Under **Device Certificate**, click **Enable**.
3. Copy the PIN ID and enter it into the **Device Certificate PIN ID** field.
4. Reenter the PIN ID into the **Confirm Device Certificate PIN ID** field.
5. Copy the PIN Value and enter it into the **Device Certificate PIN Value** field.
6. Reenter the PIN Value into the **Confirm Device Certificate PIN Value** field.

STEP 9 | Click **OK** to save the service definition.

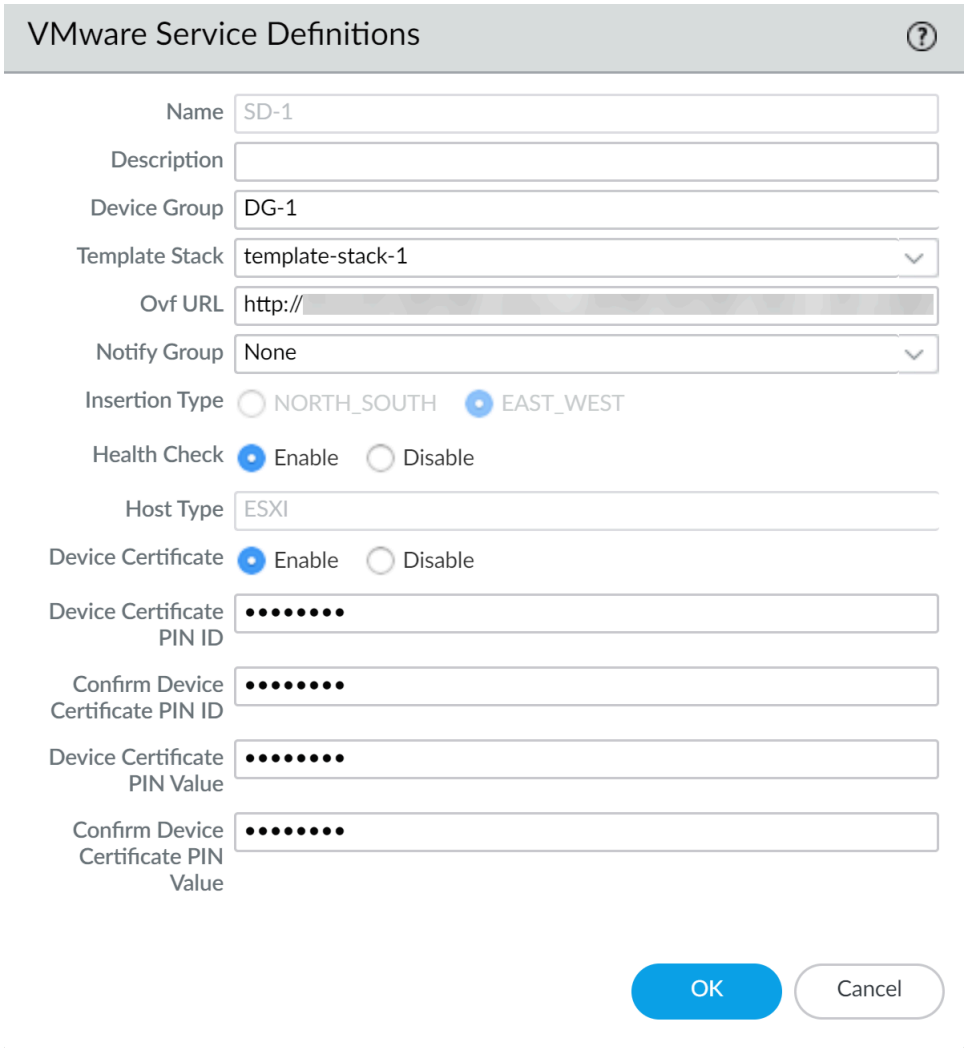

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Attach the service definition to the service manager.

You cannot use a service definition in more than one service manager.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Manager** and click the link of the service manager name.
2. Under Service Definitions, click **Add** and select your service definition from the drop-down.
3. Click **OK**.

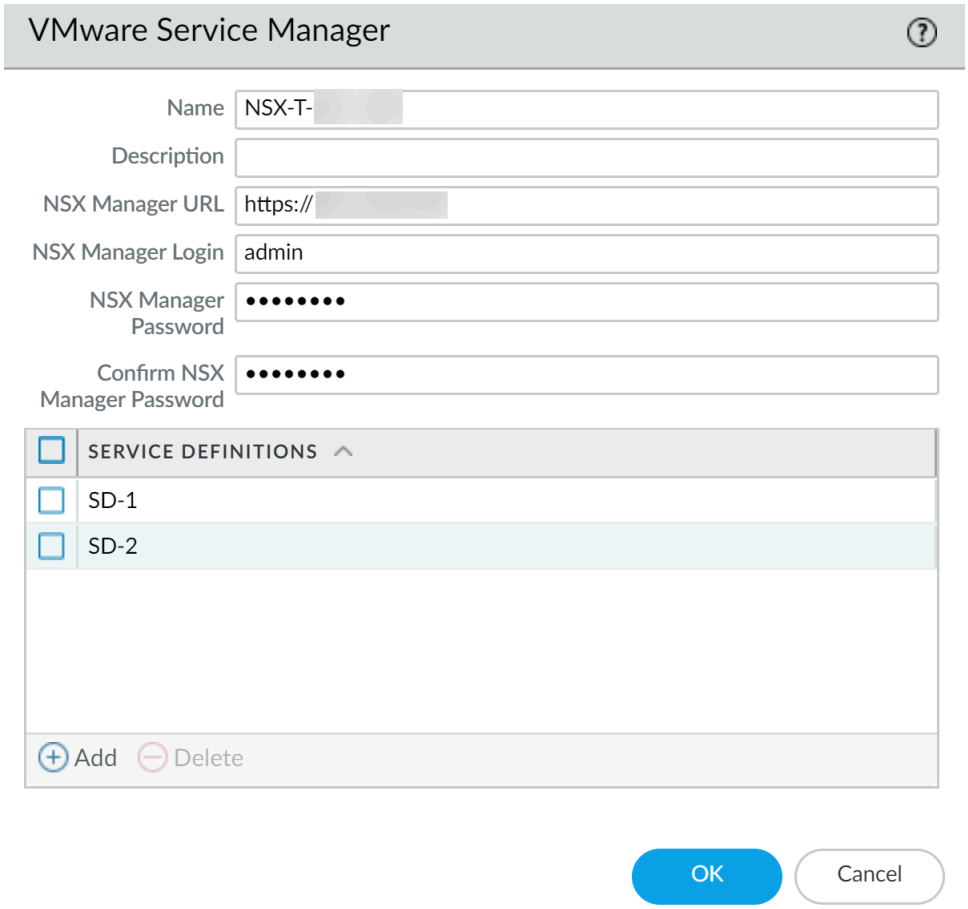

 ©2024 Palo Alto Networks, Inc.

STEP 11 | Add the authorization code to license the firewalls.

1. Select **Panorama** > **Device Groups** and choose the device group you associated with the service definition you just created.
2. Under **Dynamically Added Device Properties**, add the authorization code you received with your order fulfillment email and, optionally, select None from the **SW Version** drop-down.

   When a new firewall is deployed on NSX-T it is automatically added to the device group, licensed using the authorization code you provided, and upgraded to the PAN-OS version you specified.

   On the support portal, you can view the total number of firewalls that you are authorized to deploy and the ratio of the number of licenses that have been used to the total number of licenses enabled by your authorization code.

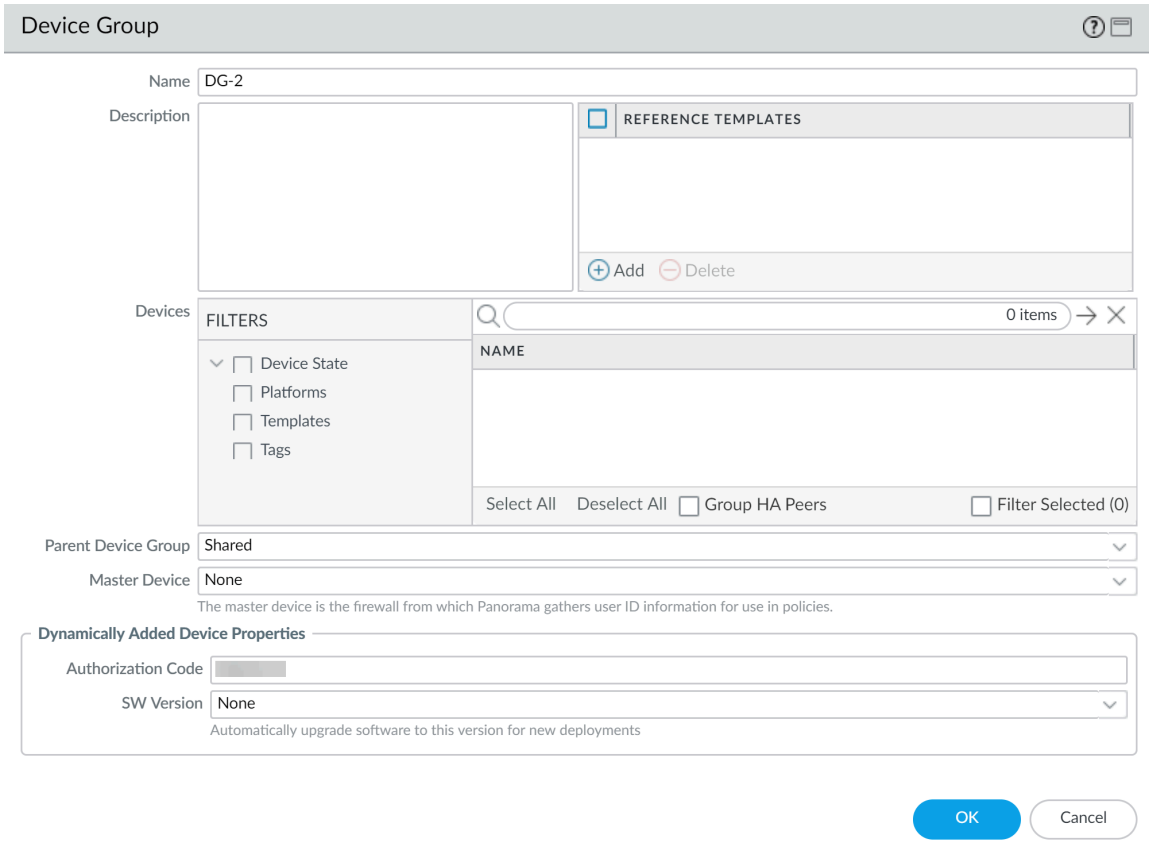

STEP 12 | **Commit to Panorama.**

 ©2024 Palo Alto Networks, Inc.

STEP 13 | On the NSX-T Manager, verify that the service definition is available.

Select **System** > **Service Deployments** > **Catalog**. The service definition is listed as a Service Instance on the NSX-T Manager.

Launch the VM-Series Firewall on NSX-T (East-West)

Complete the following procedure to deploy the VM-Series firewall as a service in your NSX-T environment. The **Deployment Specification** and **Deployment Template** fields are automatically populated with information pushed from Panorama as part of the service definition.

Do not edit any settings under Deployment Attributes. These values are imported from Panorama and changing them causes the deployment to fail.

STEP 1 | Log in to the NSX-T Manager.

STEP 2 | Select **System** > **Service Deployments** > **Deployment**.

STEP 3 | Select your service definition from the **Partner Service** drop-down.

STEP 4 | Click **Deploy Service**.

STEP 5 | Enter a descriptive **Name** for your service deployment.

STEP 6 | Select the **Compute Manager** (vCenter).

STEP 7 | Select a **Deployment Type**—**Clustered** or **Host Based**.

STEP 8 | If you selected **Clustered** as the **Deployment Type**, enter the **Clustered Deployment Count** to specify the number of VM-Series firewall instances to deploy on the cluster.

STEP 9 | Select a **Host** if you are launching the VM-Series in a clustered deployment. Select a particular host from the **Host** drop-down or **Any** to allow NSX-T Manager to choose the host. This option is grayed out in **Host Based** deployments.

STEP 10 | Select a **Data Store** as the repository for the VM-Series firewall. In a clustered deployment, select a shared data store if you choose **Any** for the host or select a local data store if you specified a particular host.

STEP 11 | Configure the **Networks** settings.

1. In the Networks column, click **Set**.
2. Select the **Network** for **eth0 - Management Nic**.
3. Select the **Network Type**—DHCP or Static IP Pool. If you choose Static IP Pool, select an **IP Pool**.
4. Check **eth1 - Data-1 Nic**.
5. Verify that both interfaces are checked.
6. Click **Save**.

 ©2024 Palo Alto Networks, Inc.

STEP 12 | Select or configure a **Service Segment**. To configure a service segment, complete the following procedure.

1. Click **Action** in the **Service Segments** column.

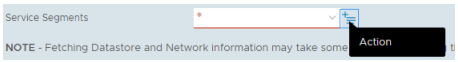

2. Click **Add Service Segment**.
3. Enter a descriptive **Name**.
4. Select a **Transport Zone (Overlay)**.
5. Click **Save** and **Close**.

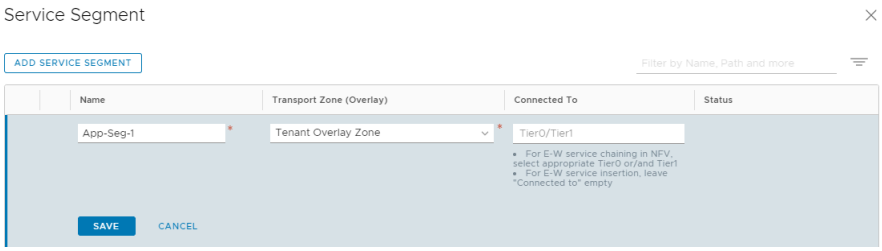

STEP 13 | Select the **Cluster** where the service will be deployed. You must select a cluster with **NSX Configuration**.

STEP 14 | Click **Save**.

STEP 15 | Verify that your firewalls deployed successfully.

1. Select **System** > **Service Deployments** > **Service Instances**.
2. Confirm that your firewalls are listed and the **Deployment Status** shows **Up**.

STEP 16 | Verify that your firewalls connected to Panorama.

1. Log in to Panorama.
2. Select **Panorama** > **Managed Devices** > **Summary**.
3. Confirm that your firewalls are listed under the correct device group and the **Device State** shows **Connected**.

   The Device Name for the VM-Series firewall is displayed on Panorama as **PA-VM:<nsx.clusterid>** for NSX-T (N-S) deployment and as **PA-VM:<nsx.servicevmid>** for NSX-T (E-W) deployment.

STEP 17 | Set a secure password for the admin account on your VM-Series firewalls.

Each VM-Series firewall uses a default username and password (admin/admin), which is used for initial login. Upon logging in for the first time, you are prompted to set a new, more secure

 ©2024 Palo Alto Networks, Inc.

password. The new password must be a minimum of eight characters and include a minimum of one lowercase and one uppercase character, as well as one number or special character.

You can update the password on each firewall individually or all at once through Panorama.

- **Panorama**—on Panorama, you can change the default password for all firewalls in a template or delete the admin user and create a new username and password.
  1. Log in to Panorama
  2. Select **Device** > **Administrators** and select the **admin** user.
  3. **Delete** the user or click the user and enter a new password.
  4. If you changed the password, click **OK**.
  5. Select **Commit** > **Push to Devices** > **Edit Selections** > **Force Template Values**.
  6. Click **OK**.
- **Firewall**—this procedure must be repeated on each VM-Series firewall.
  1. Log in to the VM-Series firewall using the default username and password.
  2. Follow the prompts to reset the password.

Create Dynamic Address Groups

A security group is a logical container that assembles guests across multiple ESXi hosts in the cluster. When you create a dynamic address group that meets the right criteria and commit your changes, a corresponding security group is created on the NSX-T Manager. Creating security groups is required to manage and secure the guests; to understand how security groups enable policy enforcement, see Policy Enforcement Using Dynamic Address Groups.

For a dynamic address group to become a security group on NSX-T, you must add match criteria in the dynamic address group in the following format: **'_nsxt_<dynamic-address-group-name>'**. The dynamic address name added in the match criteria must match the dynamic address group name exactly. For example, a dynamic address group called **applications** must include match criteria **'_nsxt_applications'**. Additionally, you must include the dynamic address group in a device group in a service definition, which is part of a service manager, and committed.

Each security group created from a dynamic address group is in the following format: **<service-def-name>_<dynamic-address-group-name>**. For example, **ServiceDef1_applications**.

Each dynamic address group you create must have a unique name across each device group configured on your Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Configure a dynamic address group for each security group required for your deployment.

1. Select **Objects** > **Address Groups**.
2. Verify that you are configuring the dynamic address groups in a device group associated with an NSX-T service definition.
3. Click **Add** and enter a **Name** and **Description** for the address group.
4. Select **Type** as **Dynamic**.
5. Define the match criteria.

   For the dynamic address group to become a security group in NSX-T Manager, the match criteria string must be enclosed in single quotes with the prefix _nsxt_ followed by the exact name of the Address Group. For example, ***'_nsxt_PAN_APP_NSX'***.
6. Repeat this process for each security group you require.

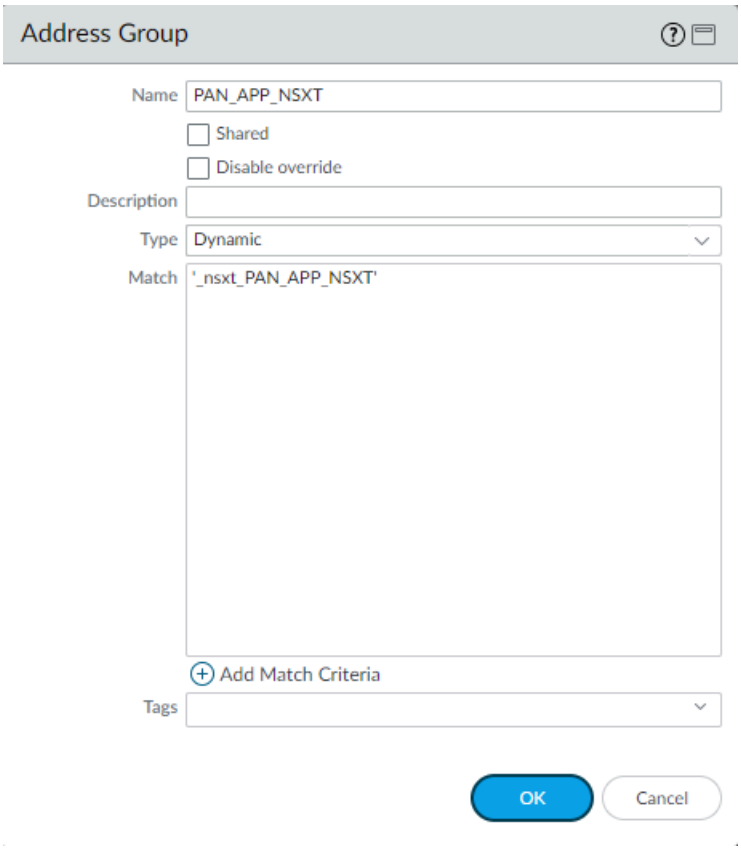

STEP 2 | **Commit** your changes.

Create Security Policies

Create security policy rules that will used to auto generate steering rules used in used in steering policy.

When you Generate Steering Rules, you will have the option to generate steering rules based on pre-rules, post-rules, or all. If you select All, the VMware plugin for NSX creates a steering rule for each applicable security in the pre and post rules. This can result in the creation of unnecessary steering rules and make managing the rules more difficult. To help easily separate your steering

 ©2024 Palo Alto Networks, Inc.

rules from your security rules, you can create your steering rules as post rules and security rules as pre rules.

To auto generate a steering rule based on a security rules created on Panorama, the security rule must meet the following criteria:

- Belongs to a parent or child device group registered with an NSX-T Service Manager.
- Is an intrazone policy and includes only one zone.
- Does not include a static address group, IP range, or netmask configured for the rule.

When deciding where to define your NSX-T steering rules in Panorama—pre or post rulesbase—consider the number of security policy rules and NSX-T steering rules you will create on Panorama and the order in which the rules are applied to traffic. Pre-rules are applied to traffic before post-rules.

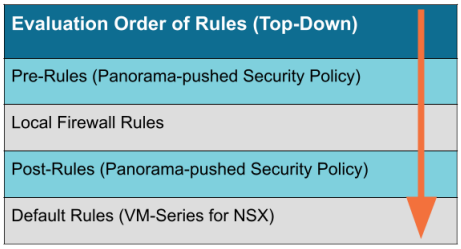

- Pre-Rules—you can use the Panorama pre-rulebase to define your NSX-T steering rules and VM-Series firewall security policy rules. If your define the security rules and steering rules in the same rulebase, you must consider the order of the security rules relative to the steering rules. When you have a large rulebase that includes both steering rules and security policy rules, it might become difficult to manage both types of rules as you scale.
- Post-Rules—separating your security policy rules used for inspection and enforcement from the security rules used to generate NSX-T steering rules can help you scale in deployment with a large amount of rules. When you auto generate your steering rules, the plugin generates a steering rule for every rule in the specified rulebase that meets the necessary criteria. Therefore, by separating the two types of rules, you can prevent unintentionally generating extraneous steering rules. Use of the post rulebase for steering rules is recommended; especially in deployments with large amounts of security policy rules.

The source and destination dynamic address groups you specify in the security rule. When you auto generate a steering rule, where the rule is applied (NSX-T Distributed Firewall or Security Group) depends on the source and destination you specified when configuring the security rule. If you selected any for the source or destination, NSX-T Manager applies the steering rule to the Distributed Firewall. If you select a dynamic address group for the source and destination, the steering is applied to the guest VMs in those security groups. If you manually create steering rules, you can specify the security group(s) where the steering rule is applied.

Ensure that your security policy that is used to define steering rules do not include dynamic address groups configured as part of an operations-centric deployment workflow. If you do, the steering rules source and destination will be pushed to NSX-T Manager as source-any and destination-any. This might impact traffic in your NSX-T environment.

If you disable a security rule that you will use to auto generate a steering rule, the steering rule will be disabled as well.

 ©2024 Palo Alto Networks, Inc.

- Use the Pre Rulebase to Define NSX-T Steering Rules
- Use the Post Rulebase to Define NSX-T Steering Rules
- Apply Security Policies to the VM-Series Firewall on NSX-T (East-West)

Use the Pre Rulebase to Define NSX-T Steering Rules

The following procedure describes how to create the security policy rules that will be used to generate NSX-T steering rules and how to create the security policy Panorama will push to the VM-Series firewaa for traffic inspection and enforcement.

Do **not** apply the traffic redirection policies unless you understand how rules work on the NSX-T Manager as well as on the VM-Series firewall and Panorama. The default policy on the VM-Series firewall is set to *deny all* traffic, which means that all traffic redirected to the VM-Series firewall will be dropped.

Create security policy rules in the associated device group. For each security rule set the Rule Type to Intrazone, select one zone in the associated template stack, and select the dynamic address groups as the source and destination. Creating a qualifying security policy in Panorama helps in the creation of a corresponding steering rule on NSX-T Manager upon steering rule generation and commit in Panorama.

STEP 1 | In Panorama, select **Policies** > **Security** > **Pre Rules**.

STEP 2 | Click **Add** and enter a **Name** and **Description** for your security policy rule.

STEP 3 | Verify that you are configuring the security rules in a device group associated with an NSX-T service definition.

STEP 4 | Set the Rule Type to **intrazone (Devices with PAN-OS 6.1 or later)**.

STEP 5 | In the Source tab, set the source zone to the zone from the template stack associated with the service definition. Then select a dynamic address group (NSX-T security group) you created previously as the Source Address. Do not add any static address groups, IP ranges, or netmasks as a Source Address.

STEP 6 | In the Destination tab, Panorama does not allow you to set a destination zone because you set the rule type to intrazone. Then select a dynamic address group (NSX-T security group) you created previously as the Destination Address. Do not add any static address groups, IP ranges, or netmasks as a Destination Address.

STEP 7 | Click **OK**.

STEP 8 | Repeat steps 1 through 7 for each steering rule you require.

STEP 9 | **Commit** your changes.

STEP 10 | Apply Security Policies to the VM-Series Firewall on NSX-T (East-West).

Use the Post Rulebase to Define NSX-T Steering Rules

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create security rules.

1. In Panorama, select **Policies** > **Security** > **Post Rules**.
2. Verify that you are configuring the security rules in a device group associated with an NSX-T service definition.
3. Click on the name of a security rule to edit.
4. Set the Rule Type to **intrazone (Devices with PAN-OS 6.1 or later)**.
5. In the Source tab, set the source zone to the zone from the template stack associated with the service definition. Then select a dynamic address group you created previously as the Source Address. Do not add any static address groups, IP ranges, or netmasks as a Source Address.
6. In the Destination tab, Panorama does not allow you to set a destination zone because you set the rule type to intrazone. Then select a dynamic address group you created previously as the Destination Address. Do not add any static address groups, IP ranges, or netmasks as a Destination Address.
7. Click **OK**.
8. Repeat steps 1 through 7 for each steering rule you require.

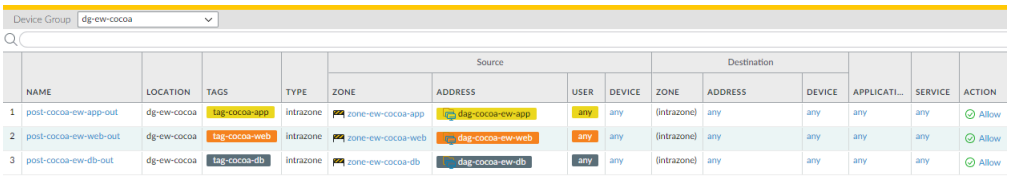

STEP 2 | **Commit** your changes to Panorama.

STEP 3 | Apply Security Policies to the VM-Series Firewall on NSX-T (East-West).

Apply Security Policies to the VM-Series Firewall on NSX-T (East-West)

Now that you have defined the steering rules, you can now use Panorama for centrally administering policies on the VM-Series firewalls.

To manage centralized policy, attach the dynamic address group as a source or destination address in security policy and push it to the firewalls; the firewalls can dynamically retrieve the IP addresses of the virtual machines that are included in each security group to enforce compliance for traffic that originates from or is destined to the virtual machines in the specified group.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create security policy rules.

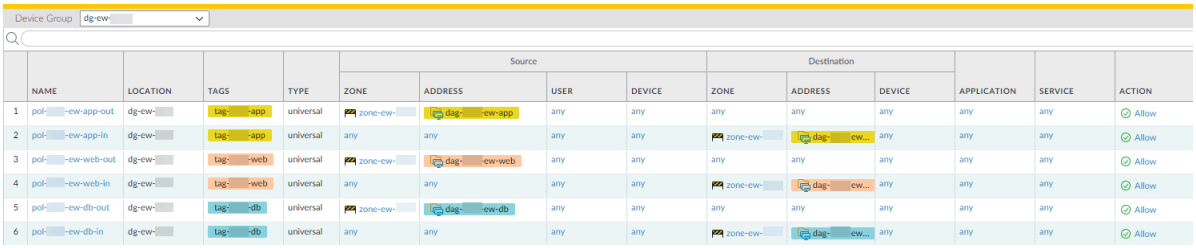

1. Select **Policies** > **Security** > **Prerules**.
2. Select the **Device Group** that you created for managing the VM-Series firewalls on NSX-T in Create Template Stacks and Device Groups on Panorama.
3. Click **Add** and enter a **Name** and a **Description** for the rule. In this example, the security rule allows all traffic between the WebFrontEnd servers and the Application servers.
4. Select the **Source Zone** and **Destination Zone**. The zone name must be the same in both columns.
5. For the **Source Address** and **Destination Address**, select or type in an address, address group or region. In this example, we select an address group, the Dynamic address group you created previously.
6. Select the **Application** to allow. In this example, we create an **Application Group** that includes a static group of specific applications that are grouped together.
   1. Click **Add** and select **New Application Group**.
   2. Click **Add** to select the application to include in the group.
   3. Click **OK** to create the application group.
7. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles for antivirus, anti-spyware, and vulnerability protection, under Profiles.
8. Repeats the steps above to create the pertinent policy rules.
9. Click **Commit**, select Commit Type as **Panorama**. Click **OK**.

STEP 2 | Apply the policies to the VM-Series firewalls for NSX-T.

1. Click **Commit**, and select Commit Type **Device Groups**.
2. Select the device group, NSX-T Device Group in this example and click **OK**.
3. Verify that the commit is successful.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Validate that the members of the dynamic address group are populated on the VM-Series firewall.

1. From Panorama, switch device context to launch the web interface of a firewall to which you pushed policies.
2. On the VM-Series firewall, select **Policies** > **Security**, and select a rule.
3. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.

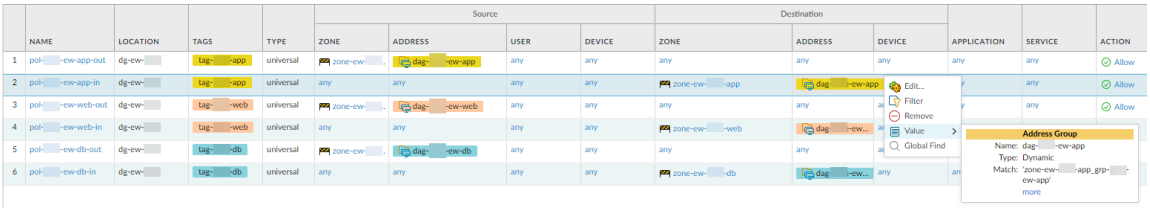

4. Click the **more** link and verify that the list of registered IP addresses is displayed.

   Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

STEP 4 | (Optional) Use template to push a base configuration for network and device configuration such as DNS server, NTP server, Syslog server, and login banner.

Refer to the Panorama Administrator's Guide for information on using templates.

STEP 5 | Create a Zone Protection profile and attach it to a zone.

A zone protection profile provides flood protection and has the ability to protect against port scanning, port sweeps and packet-based attacks. It allows you to secure intra-tier and inter-tier traffic between virtual machines within your data center and traffic from the Internet that is destined to the virtual machines (workloads) in your data center.

1. Select your **Template**.
2. Select **Network** > **Network Profiles** > **Zone Protection** to add and configure a new profile.
3. Select **Network** > **Zones**, click the default-zone listed and select the profile in the **Zone Protection Profile** drop down.

STEP 6 | Create a DoS Protection profile and attach it to DoS Protection policy rule.

1. Select your **Device Group**.
2. Select **Objects** > **Security Profiles** > **DoS Protection** to add and configure a new profile.
   - A classified profile allows the creation of a threshold that applies to a single source IP. For example, you can configure a max session rate for an IP address that matched the policy, and then block that single IP address once the threshold is triggered.
   - An aggregate profile allows the creation of a max session rate for all packets matching the policy. The threshold applies to new session rate for all IP addresses combined. Once the threshold is triggered it affects all traffic that matches the policy.
3. Create a new DoS Protection policy rule in **Policy** > **DoS Protection**, and attach the new profile to it.

 ©2024 Palo Alto Networks, Inc.

Create Dynamic Address Group Membership Criteria

In NSX-T, you can configure the membership criteria for your virtual machines and IP set belonging to an NSX-T security group (dynamic address group) in the Panorama plugin for NSX. For each dynamic address group, you must specify a service definition and define up to five match criteria and each criterion includes up to five match rules.

You create this membership criteria on the plugin and then push it to NSX-T Manager. However, this does not apply the membership criteria to guest virtual machines in your deployment. You must define and apply membership data, such as tags, to your guest VMs in NSX-T Manager.

The rules that the Panorama plugin for NSX-T identifies and classifies virtual machines based on two membership types—Virtual Machine or IP set. The keys and operators usable with each member type are listed in the table below.

Member Type	Key	Operator
IP Set	Tag	Equals
Virtual Machine	• Tag • Name • OS Name • Computer Name	• Equals • Contains • Starts With • Ends With • Not Equals (Not applicable with Tag key)

*Membership criteria changes should be made only on Panorama; do not make changes on NSX-T Manager. If you make changes on NSX-T Manager, the Panorama plugin for VMware NSX show the service definition as out-of-sync. You should click on the **Out-of-Sync** link to see the specific reason for the out-of-sync status. If a membership criteria change is the cause, perform a configuration sync by clicking **NSX-T Config-Sync**.*

STEP 1 | Select **Panorama** > **VMware** > **NSX-T** > **Membership Criteria** > **Add**.

To add or modify membership criteria for a service definition, with at least one dynamic address group, you can click on the service definition name instead of clicking **Add**.

STEP 2 | From the **Name**, select a service definition for the Membership Criteria. The selected service definition must have East_West insertion type and used as part of a security-centric deployment.

STEP 3 | Click **Add** to specify a dynamic address group.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Select a **Dynamic Address Group** from the drop-down. The drop-down lists the dynamic address groups associated with the specified service definition.

> *The plugin UI displays dynamic and static address groups configured on Panorama. Take care not accidently select a static address group when configuring membership criteria.*

STEP 5 | Click **Add** to define the criteria associated with the chosen dynamic address group.

STEP 6 | Enter a descriptive name for the **Criteria**.

STEP 7 | Click **Add** to define a rule.

STEP 8 | Define a rule. You can create up to five rules.

1. Enter a descriptive name for the rule.
2. Select the **Member Type**—Virtual Machine or IP Set.
3. Select the **Key**—Tag, Name, OS Name, Computer Name.
4. Select the **Operator**—Equals, Contains, Starts With, Ends With, Not Equals.
5. Enter the **Value**.

   If the Key is set to Tag, the Value is the Tag. The plugin user interface does not list the Tags, so you must use the Panorama CLI (with NSX-T Manager 3.0.0. and later).

   `request plugins vmware_nsx nsx_t nsxt-tags service-definition <SD_name>`
6. (Optional) Enter the **Scope**. Scope is applicable only with the key **Tag**. Scope is an optional value applied to an object tag in NSX-T. The scope is defined on NSX-T

 ©2024 Palo Alto Networks, Inc.

Manager. For example, if you tag virtual machines based on operating system, you can create tags for Windows, Linux, and MacOS and then set the scope of each tag to OS.

To view the tags and scope, use the Panorama CLI (with NSX-T Manager 3.0.0 and later).

Execute the following command to view the list of tags.

```
request plugins vmware_nsx nsx_t nsxt-tags service-definition <SD_name>
```

Execute the following command to view the scope associated with the specified tag.

```
request plugins vmware_nsx nsx_t nsxt-scope tag <tag_value> service-definition <SD-name>
```

7. Click **OK**.
8. (Optional) Click **Add** to create additional (up to five total) rules.

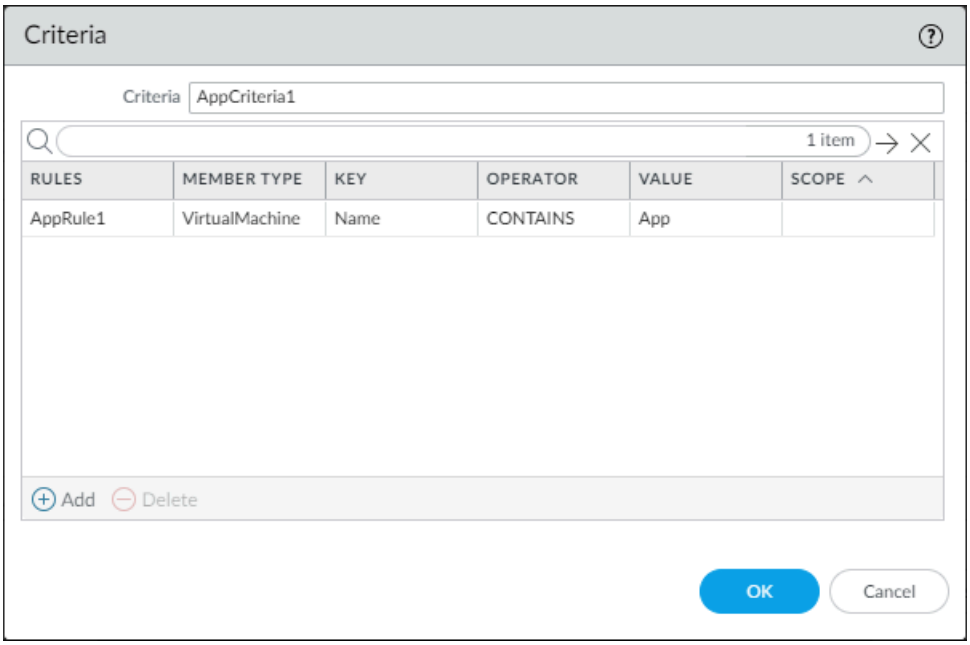

 ©2024 Palo Alto Networks, Inc.

STEP 9 | On the Dynamic Address Group window, click **OK** to finish or **Add** to create additional criteria (up to five total) and rules.

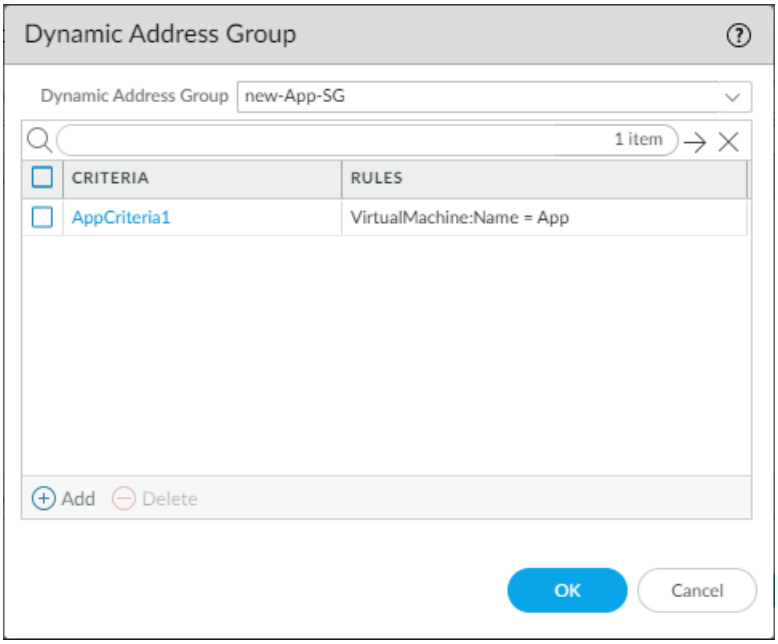

STEP 10 | On the Membership Criteria window, click **OK** to finish or **Add** to specify additional dynamic address groups.

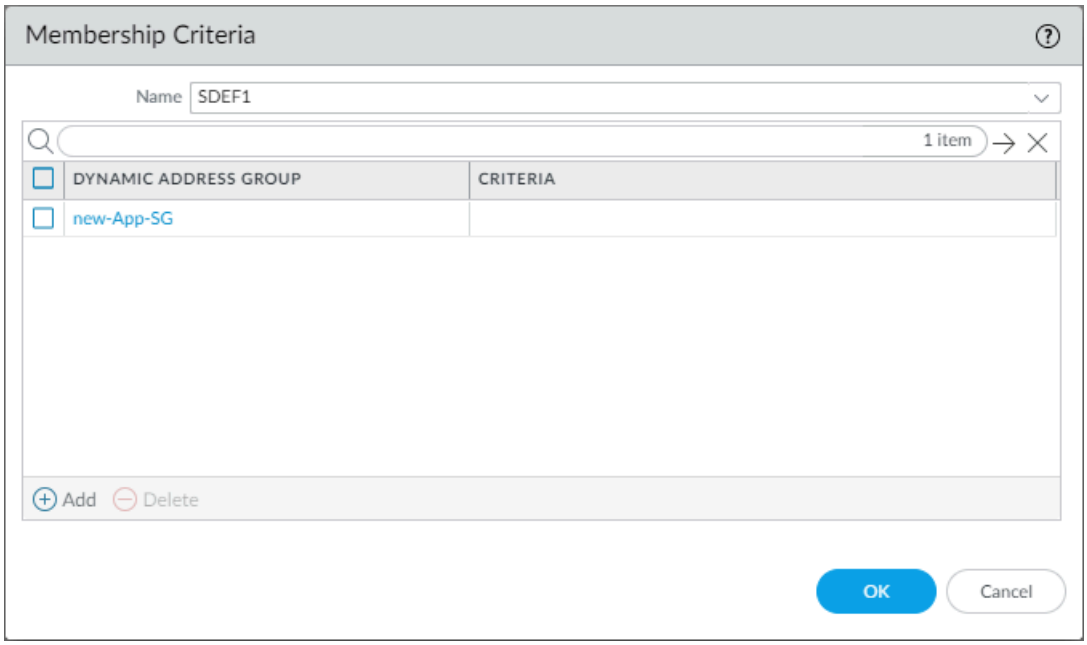

Generate Steering Policy

Steering policy is used by NSX-T to define the service chain to which traffic will be steered. You can create steering policy manually or you can auto generate steering policy.

When you auto generate steering policy, the Panorama plugin for VMware NSX-T creates a steering policy for each specified service manager and the associated service definitions. By

 ©2024 Palo Alto Networks, Inc.

default, TCP strict is disabled and the Failure Policy is set to Allow. Auto-generated policy uses the **auto_<service-def-name>_<zone-name>_steering_policy** naming format.

When TCP Strict is enabled, the firewall enforces the requirement of the three-way handshake. If the firewall picks up traffic mid-session (for example, due to asymmetric traffic) and does not detect a three-way handshake, the session is dropped. See VMware NSX-T documentation for more information.

The Failure Policy defines what happens to traffic if the firewall goes down. If you select Allow, the traffic continues on to its destination. If you select Block, the traffic is dropped.

Additionally, you have the option to select all your service managers instead of selecting specific service managers. Choosing **All** is not recommended if any of your service managers contain operations-centric service definitions. The plugin will create steering policy for each zone associated with the operation-centric service definitions and then push it to NSX-T Manager. If you do choose **All**, verify that the service manager you select when you auto generating steering policy includes only security-centric service definitions.

If you auto-generate steering policy, you must also auto-generate steering rules. And you manually create steering policy, you must also manually create steering rules.

- Auto Generate Steering Policy
- Manually Create Steering Policy

*Steering policy changes should be made only on Panorama; do not make changes on NSX-T Manager. If you make changes on NSX-T Manager, the Panorama plugin for VMware NSX show the service definition as out-of-sync. You should click on the **Out-of-Sync** link to see the specific reason for the out-of-sync status. If a steering policy change is the cause, perform a configuration sync by clicking **NSX-T Config-Sync**.*

Auto Generate Steering Policy

Use the following procedure to auto generate steering policy.

*The following steps are for specifying service managers instead of selecting **All**.*

STEP 1 | Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Policy**.

STEP 2 | Click **Auto Generate**.

STEP 3 | For **Service Managers**, choose **Select**.

*If you select **All** instead of selecting specific service managers, the plugin will generate steering policy for each service definition associated with each service manager in your configuration. Additionally, make sure that your selected service manager includes security-centric service definitions.*

STEP 4 | Click **Add** to select the service manager.

STEP 5 | Select a **Service Manager** from the drop-down.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Click **Add** to select the service definitions.

STEP 7 | Select the service definition from the drop-down.

STEP 8 | Click **OK** and click **OK** again.

STEP 9 | **Commit** your changes to Panorama.

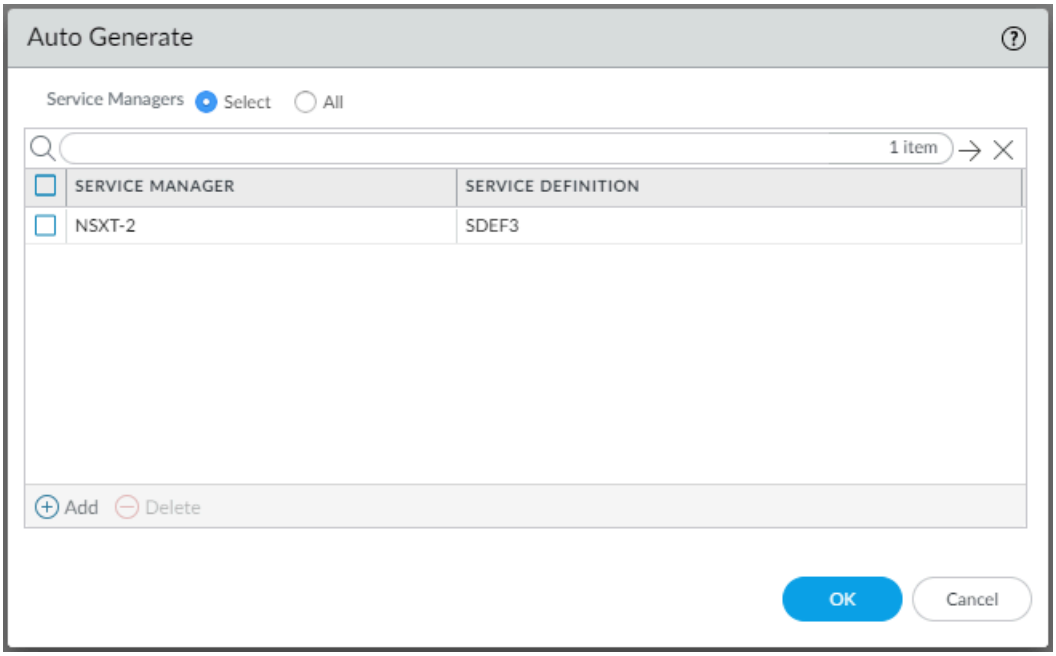

Manually Create Steering Policy

Use the following procedure to manually create steering policy.

STEP 1 | Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Policy**.

STEP 2 | Click **Add**.

STEP 3 | Enter a descriptive **Name** for your steering policy.

The steering policy name cannot include any spaces.

STEP 4 | Select a **Service Definition** from the drop-down.

STEP 5 | Select a **Service Chain** from the drop-down.

STEP 6 | (Optional) Enable **TCP Strict**. This option is disabled by default.

STEP 7 | Choose the **Failure Policy**— **Allow** or **Block**. Allow is the default.

STEP 8 | Click **OK**.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | **Commit** your changes to Panorama.

Generate Steering Rules

Steering rules are defined in steering policy. A rule defines the source and destination of the traffic, introspection services, the NSX-T objects the rule is applied to, and the traffic redirection policy. You can create steering rules manually or generate steering rules automatically.

You must generate or create steering policy before generating or creating steering rules.

To auto generate a steering rule based on a security rules created on Panorama, the security rule must meet the following criteria:

- Belongs to a parent or child device group registered with an NSX-T Service Manager.
- Is an intrazone policy and includes only one zone.
- Does not include a static address group, IP range, or netmask configured for the rule.

Auto-generated steering rules uses the **auto_<device-group-name>_<device-group-rule-name>** naming format.

By default, auto-generated steering rules are configured without an NSX services specified. Additionally, the NSX Traffic Direction is set to in-out, Logging is disabled, IP protocol is ipv4-ipv6, and the Action is set to redirect. After auto-generating rules, you can update the steering to change the default values.

Additionally, you have the option to select all your service managers instead of selecting specific service managers. Choosing **All** is not recommended.

If you auto-generate steering policy, you must also auto-generate steering rules. And if you manually create steering policy, you must also manually create steering rules.

- Auto Generate Steering Rules
- Manually Create Steering Rules

*Steering rules changes should be made only on Panorama; do not make changes on NSX-T Manager. If you make changes on NSX-T Manager, the Panorama plugin for VMware NSX show the service definition as out-of-sync. You should click on the **Out-of-Sync** link to see the specific reason for the out-of-sync status. If a steering rules change is the cause, perform a configuration sync by clicking **NSX-T Config-Sync**.*

 ©2024 Palo Alto Networks, Inc.

Auto Generate Steering Rules

Use the following procedure to auto generate steering rules.

When you auto generate a steering rule, where the rule is applied (NSX-T Distributed Firewall or Security Group) depends on the source and destination you specified when configuring the security rule. If you selected **Any** for the source or destination, NSX-T Manager applies the steering rule to the Distributed Firewall. If you select a dynamic address group for the source and destination, the steering is applied to the guest VMs in those security groups.

If you make any changes to device group configuration that is also part of steering rule configuration, such as source and destination address group that map to the Applied To setting in a steering rule, you must auto generate the steering rule again for the changes to take effect.

*The following steps are for specifying service managers instead of selecting **All**.*

STEP 1 | Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Rule**.

STEP 2 | Click **Auto Generate**.

STEP 3 | Select the type of Security Rules from the drop-down—**All**, **Pre Rulebase** only, or **Post Rulebase** only. The security rules are pulled from the service definitions specified in the following steps.

If you regenerate steering rules, all current rules are deleted and new rules are created based on the selected rule base. If you originally created steering rules using the Pre Rulebase and then regenerate steering rules using the Post Rulebase, only the post-rulebase steering rules will remain.

STEP 4 | For **Type**, choose **Select**.

STEP 5 | Click **Add** to specify the **Service Manager**(s) and **Service Definition**(s).

STEP 6 | Select a **Service Manager** from the drop-down.

STEP 7 | Click **Add** to select the service definition(s).

STEP 8 | Click **OK**.

STEP 9 | Click **OK** to finish or **Add** to specify additional service managers and service definitions.

STEP 10 | (Optional) Click on an auto-generated rule to modify the following default options.

If you regenerate steering rules, any changes you made to a previously-generate steering rule will be overwritten.

- Enable NSX-T **Logging**.
- Click **Add** to specify **NSX Services**, such as Active Directory Server, HTTPS, DNS, etc.
- **Disable** the rule. If you disable a steering rule but the corresponding security rule is enabled (**Device Group** > **Policies** > **Security**), the steering rule will also be enabled.
- **Applied to** allows you change where the steering rule is applied—**DFW** or **Security Group**.

 ©2024 Palo Alto Networks, Inc.

STEP 11 | Clean up unwanted or incorrect steering rules.

If, for example, your device group contains security rules in the same rulebase as your NSX-T steering rules, the plugin generates security rules based on those non-NSX-T security rules. Because those rules do not refer to an NSX-T dynamic address group, the source and destination for those rules will be set to Any Any in NSX-T Manager. This condition can impact how NSX-T Manager directs traffic. To avoid this, you must manually delete the incorrect steering rules.

1. Select the incorrect steering rules.
2. Click **Delete**.
3. Click **Yes** to confirm the deletion.

STEP 12 | **Commit** your configuration to push it to NSX-T Manager.

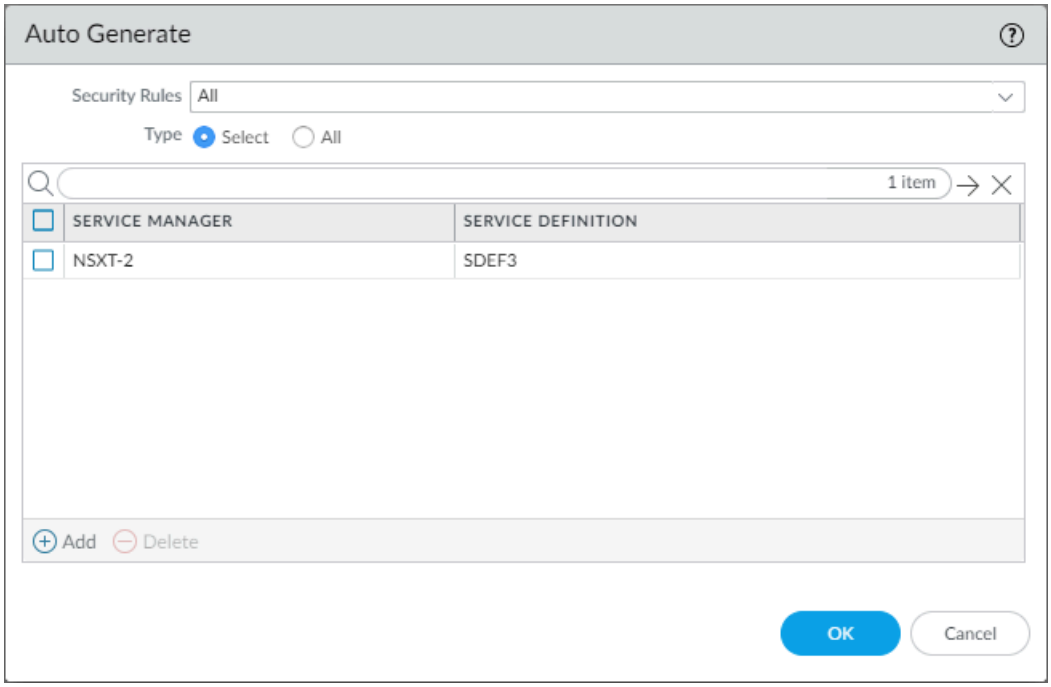

Manually Create Steering Rules

Use the following procedure to manually create steering rules.

STEP 1 | Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Rule**.

STEP 2 | Click **Add**.

STEP 3 | Enter a descriptive **Name** for the steering rule.

The steering rule name cannot include any spaces.

STEP 4 | Select a **Steering Policy** from the drop-down.

STEP 5 | Select a **Device Group** from the drop-down.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Select a **Security Rule** from the drop-down.

> *The Security Rule drop-down displays rules from all security rules across all device groups of Service Definition. Ensure you select the appropriate security rule.*

STEP 7 | Specify the **Action**—**Redirect** or **Do Not Redirect**.

STEP 8 | (Optional) Enable NSX-T **Logging**.

STEP 9 | Specify the **IP Protocol**—**ipv4-ipv6**, **ipv4**, or **ipv6**.

STEP 10 | Specify the **NSX Traffic Direction**—**in-out**, **in**, or **out**.

STEP 11 | (Optional) Click Add to specify **NSX Services**, such as Active Directory Server, HTTPS, DNS, etc.

> *The following ALG services are not supported: FTP, TFTP, ORACLE_TNS, SUN_RPC_TCP, SUN_RPC_UDP, MS_RPC_TCP, MS_RPC_UDP, NBNS_BROADCAST, NBDG_BROADCAST.*

STEP 12 | **Applied To**—**DFW** or **Security Groups**. You can select one or more security group. Security groups are created from dynamic address groups configured on Panorama. The security group names are formatted as follows **<servicedefinition>_<dynamic-address-group>**. If you select DFW, the steering rule is applied to all guest VMs, regardless of their security membership.

STEP 13 | (Optional) Disable the rule.

STEP 14 | Click **OK**.

 ©2024 Palo Alto Networks, Inc.

STEP 15 | **Commit** your configuration to push it to NSX-T Manager.

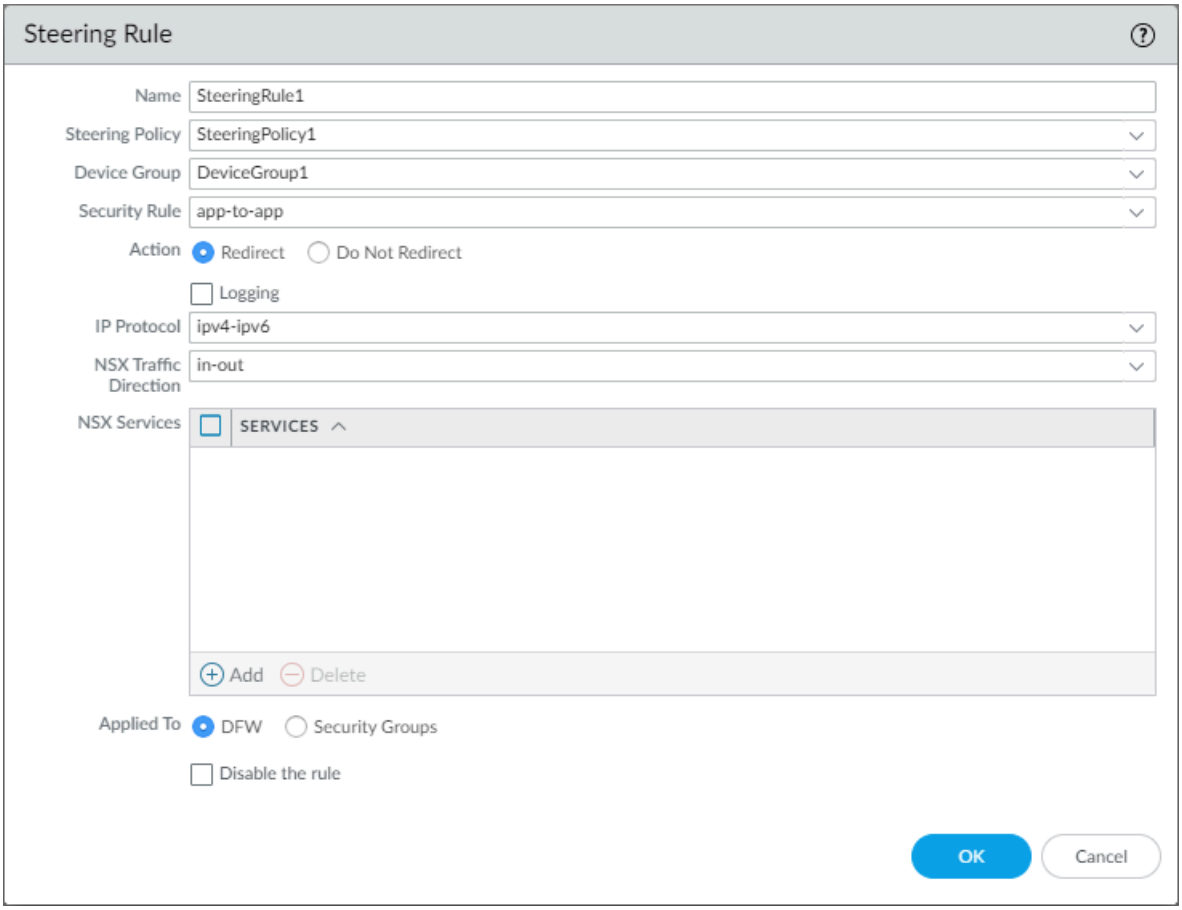

Delete a Service Definition from Panorama

Complete the following procedure to delete a service definition from your NSX-T configuration on Panorama.

STEP 1 | Log in to Panorama.

STEP 2 | Delete the steering rules and steering policy associated with the service definition to be deleted.

1. Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Rules**.
2. Select the steering rules to be deleted.
3. Click **Delete**.
4. Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Policy**.
5. Select the steering policy to be deleted.
6. Click **Delete**.

STEP 3 | **Commit** your changes.

STEP 4 | Delete the VM-Series firewalls deployed in NSX-T that are associated with service definition to be deleted.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Delete the membership criteria associated with the service definition.

1. Select **Panorama** > **VMware** > **NSX-T** > **Membership Criteria**.
2. Select the criteria to be deleted.
3. Click **Delete**.

STEP 6 | Unlink the service definition from the service manager with which it is associated.

1. Select **Panorama** > **VMware** > **NSX-T** > **Service Managers**.
2. Click the service manager name.
3. Select the service definition.
4. Click **Delete**.
5. Click **OK**.

STEP 7 | **Commit** your changes.

Migrate from VM-Series on NSX-T Operation to Security Centric Deployment

Use the following procedure to migrate your operations-centric NSX-T deployment to a security-centric NSX-T deployment.

STEP 1 | Log in to Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Modify the match criteria of your dynamic address groups to follow the format required for a security-centric deployment.

1. Select **Objects** > **Address Groups**.
2. Verify that you are configuring the dynamic address groups in a device group associated with an NSX-T service definition.
3. Click on the name of a previously created NSX-T dynamic address group.
4. Edit the match criteria.

   For the dynamic address group to become a security group in NSX-T Manager, the match criteria string must be enclosed in single quotes with the prefix _nsxt_ followed by the exact name of the Address Group. For example, ***'_nsxt_PAN_APP_NSX'***.
5. Repeat this process for each security group you require.

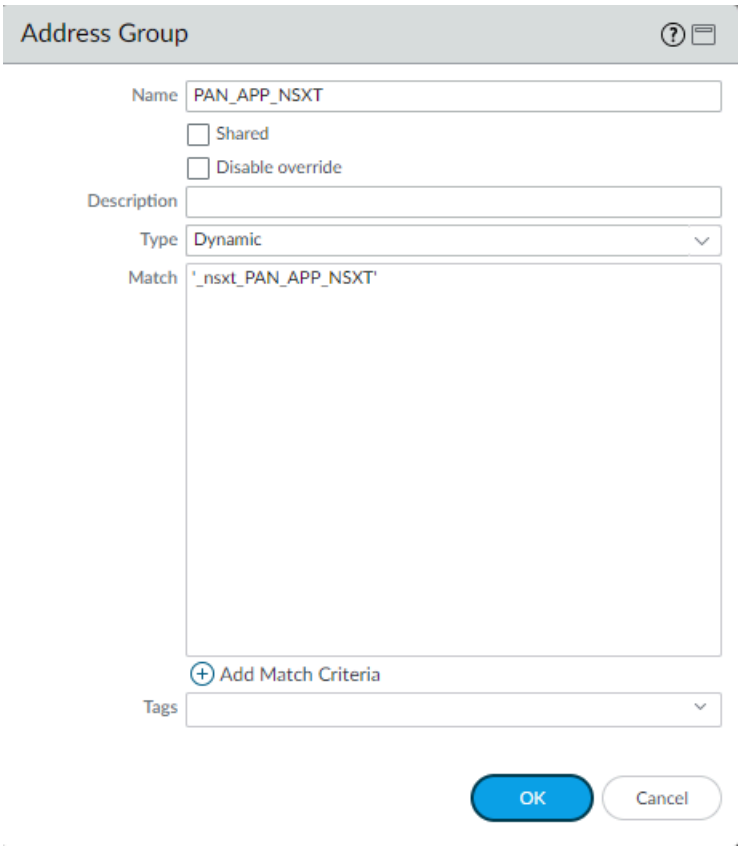

STEP 3 | Set the security rules to be as NSX-T steering rules to intrazone.

1. In Panorama, select **Policies** > **Security** > **Pre Rules**.
2. Verify that you are configuring the security rules in a device group associated with an NSX-T service definition.
3. Click **Add** and enter a **Name** and **Description** for your security policy rule.
4. Set the Rule Type to **intrazone (Devices with PAN-OS 6.1 or later)**.
5. In the Source tab, set the source zone to the zone from the template stack associated with the service definition. Then select a dynamic address group you created previously

 ©2024 Palo Alto Networks, Inc.

as the Source Address. Do not add any static address groups, IP ranges, or netmasks as a Source Address.

6. In the Destination tab, Panorama does not allow you to set a destination zone because you set the rule type to intrazone. Then select a dynamic address group you created previously as the Destination Address. Do not add any static address groups, IP ranges, or netmasks as a Destination Address.
7. Click **OK**.
8. Repeat steps 1 through 7 for each steering rule you require.
9. **Commit** your changes.

STEP 4 | Auto generate new steering policy.

The following steps are for specifying service managers instead of selecting **All**.

1. Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Policy**.
2. Click **Auto Generate**.
3. For **Service Managers**, choose **Select**.

   *If you select **All** instead of selecting specific service managers, the plugin will generate steering policy for each service definition associated with each service manager in your configuration.*
4. Click **Add** to select the service manager.
5. Select a **Service Manager** from the drop-down.
6. Click **Add** to select the service definitions.
7. Select the service definition from the drop-down.
8. Click **OK** and click **OK** again.

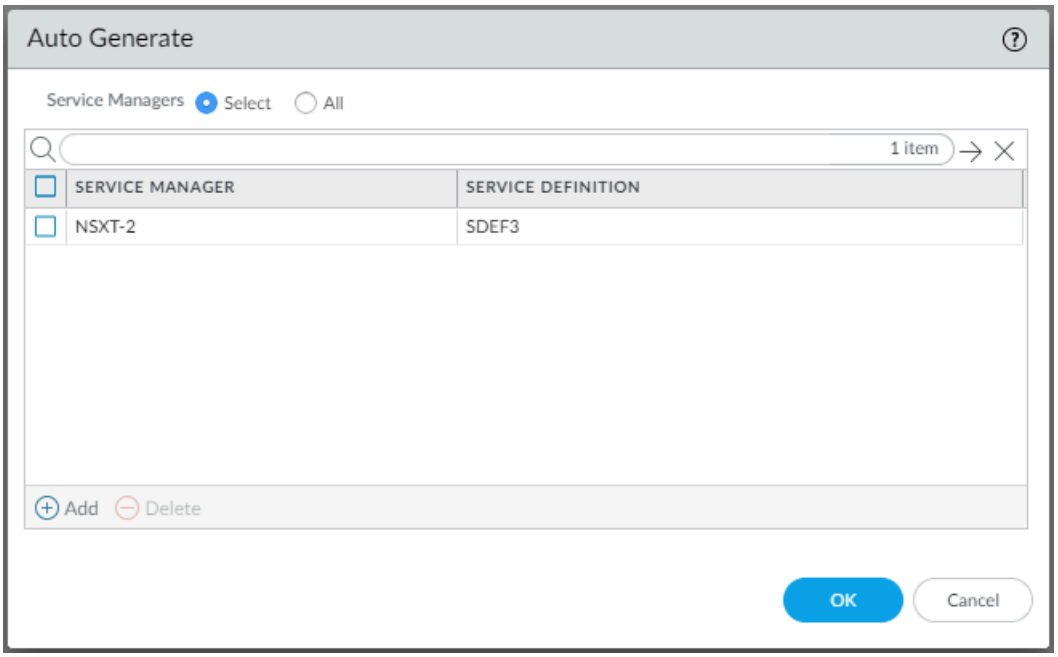

9. **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Auto generate new steering rules.

If you auto-generate steering policy, you must also auto-generate steering rules. And if you manually create steering policy, you must also manually create steering rules.

*The following steps are for specifying service managers instead of selecting **All**.*

1. Select **Panorama** > **VMware** > **NSX-T** > **Network Introspection** > **Rule**.
2. Click **Auto Generate**.
3. Select the type of Security Rules from the drop-down—**All**, **Pre Rulebase** only, or **Post Rulebase** only. The security rules are pulled from the service definitions specified in the following steps.
4. For **Type**, choose **Select**.
5. Click Add to specify the **Service Manager**(s) and **Service Definition**(s).
6. Select a **Service Manager** from the drop-down.
7. Click **Add** to select the service definition(s).
8. Click **OK**.
9. Click **OK** to finish or **Add** to specify additional service managers and service definitions.
10. (Optional) Click on an auto-generated rule to modify the default options.

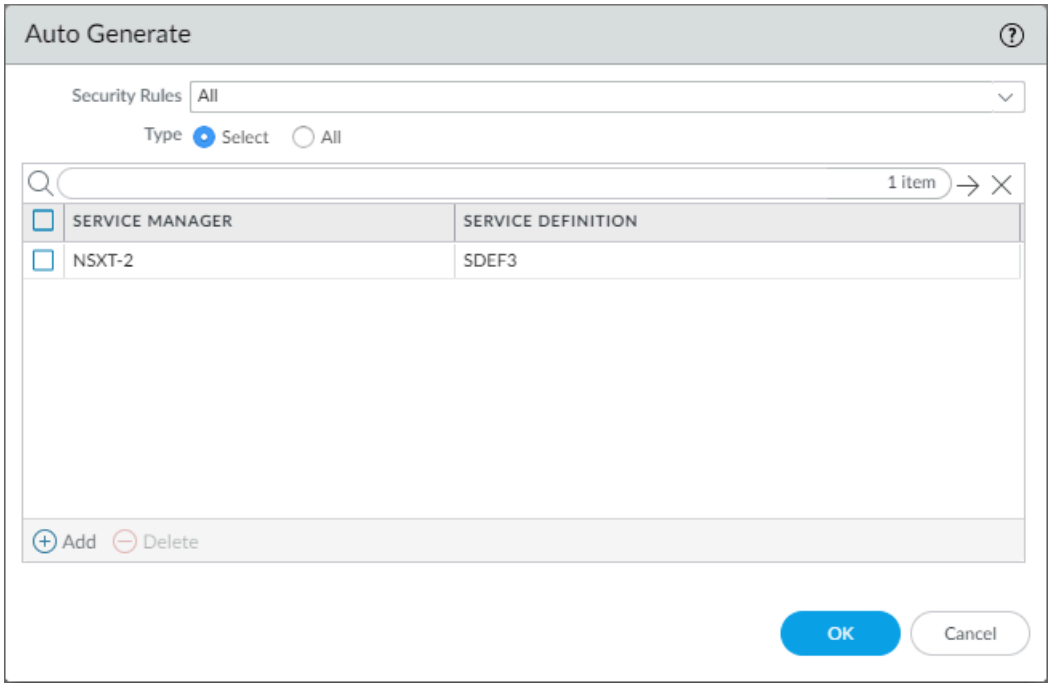

STEP 6 | Create Dynamic Address Group Membership Criteria.

STEP 7 | **Commit** your changes to Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Delete the operations-centric steering rules from NSX-T Manager.

1. Log in to NSX-T Manager.
2. Select **Security** > **Network Introspection (E-W)** > **Rules**.
3. Select each operations-centric steering rules.
4. Click **Delete**.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Delete the operations-centric service chain from NSX-T Manager.

1. Log in to NSX-T Manager.
2. Select **Security** > **Network Introspection Settings** > **Service Chains**.
3. Click the vertical ellipses.
4. Click **Delete**.

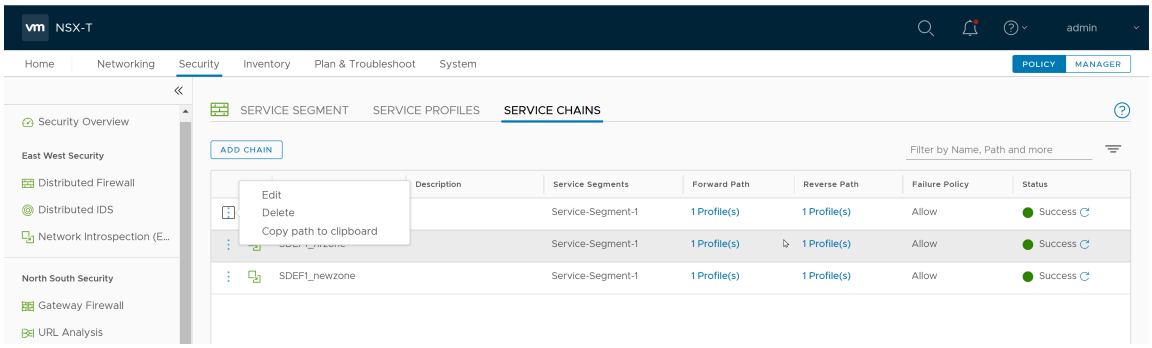

 ©2024 Palo Alto Networks, Inc.

Extend Security Policy from NSX-V to NSX-T

If you are moving from an NSX-V deployment to an NSX-T deployment or combining an NSX-T deployment with an NSX-V deployment, you can extend your existing security policy from NSX-V to NSX-T without having to recreate the policy rules. This is achieved by leveraging your existing device groups and sharing them between the NSX-V and NSX-T service definitions. After migrating your policy to NSX-T, you can continue using the VM-Series for NSX-V or remove your NSX-V deployment.

STEP 1 | Install the Panorama Plugin for VMware NSX 3.2.0 or later. See the Panorama Plugin for VMware NSX 3.2.0 Release Notes before upgrading.

STEP 2 | Configure an NSX-T service definition for each NSX-V service definition in your deployment. Do not create new device groups; instead use your existing NSX-V device groups. Using the existing device groups allows you to apply the same security policy rules used on NSX-V to the VM-Series firewalls deployed on NSX-T. If you have policy that reference a particular zone, add the same template stack from your NSX-V service definition to your NSX-T service definition. Additionally, if your device group references a particular template, ensure that you select the template stack that includes the template referenced in the device group.

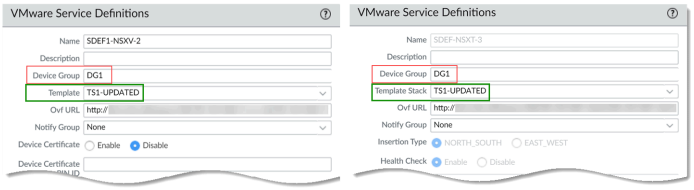

STEP 3 | Configure an NSX-T service manager and associate the NSX-T service definitions to the service manager.

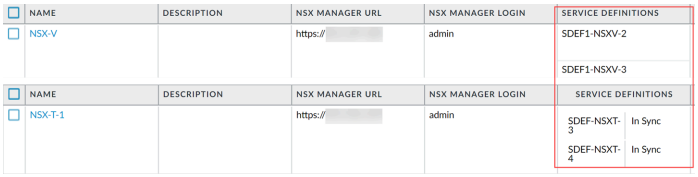

STEP 4 | Prepare your NSX-T environment and deploy the VM-Series firewall. You must create your security groups, service chains, and traffic redirection policy before launching the VM-Series firewall.

- Deploy the VM-Series Firewall on NSX-T (North-South)
- Deploy the VM-Series Using the Operations-Centric Workflow

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Add the NSX-T tags to you existing dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click on the name of an existing NSX-V dynamic address group.
3. Click **Add Match Criteria** to display the tags from NSX-V and NSX-T.
4. Add the NSX-T tag to the dynamic address groups. Be sure to use the **OR** operator between the tags.
5. When you have added all the necessary tags, click **OK**.
6. **Commit** your changes.

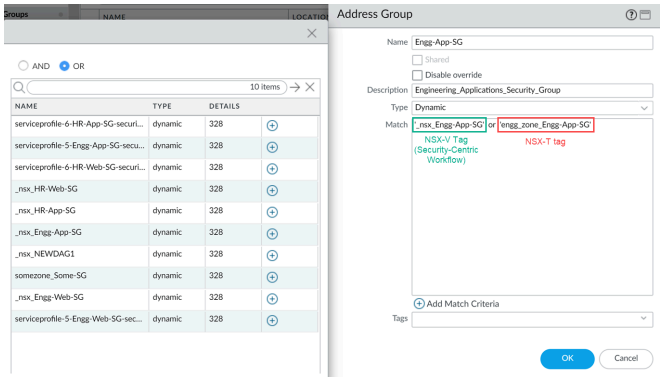

STEP 6 | After your VM workloads have successfully migrated from NSX-V to NSX-T, you remove the NSX-V tags from your dynamic address groups if you plan to discontinue use of NSX-V. All NSX-V tags and corresponding IP addresses are unregistered after all NSX-V related configuration is removed from the Panorama plugin for NSX and VM-Series firewall configuration is removed from NSX-V manager.

Use In-Place Migration to Move Your VM-Series from NSX-V to NSX-T

Complete the following procedure to migrate your VM-Series firewall configuration from NSX-V to NSX-T. By migrating your configuration, you can reuse policy and dynamic address groups already configured on Panorama. This procedure refers to information and processes published in VMware documentation as well as steps specific to PAN.

This procedure supports operations-centric NSX-V deployments only. An operations-centric deployment means that your policy rules for redirecting traffic to the VM-Series firewall were created in NSX-V Manager, not Panorama.

This procedure requires NSX-T Manager 3.1.0 or later.

It is recommended that plan for security downtime while performing this migration.

STEP 1 | Prepare your NSX-V and NSX-T environments for migration based on the steps described by VMware.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Install the Panorama Plugin for VMware NSX 3.2.0 or later. See the Panorama Plugin for VMware NSX 3.2.0 Release Notes before upgrading.

STEP 3 | Enable Communication Between NSX-T Manager and Panorama.

STEP 4 | Configure an NSX-T service definition for each NSX-V service definition in your deployment. **Do not create new device groups**; instead use your existing NSX-V device groups. Using the existing device groups allows you to apply the same security policy rules used on NSX-V to the VM-Series firewalls deployed on NSX-T. If you have policy that reference a particular zone, add the same template stack from your NSX-V service definition to your NSX-T service definition. Additionally, if your device group references a particular template, ensure that you select the template stack that includes the template referenced in the device group.

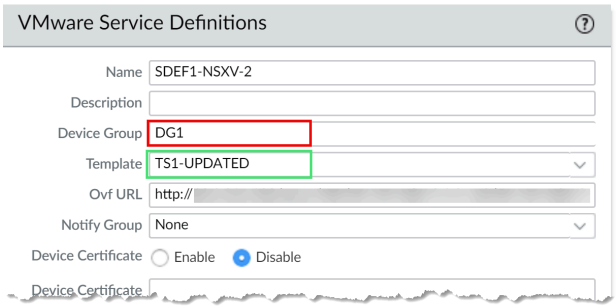

STEP 5 | Configure an NSX-T service manager and associate the NSX-T service definitions to the service manager.

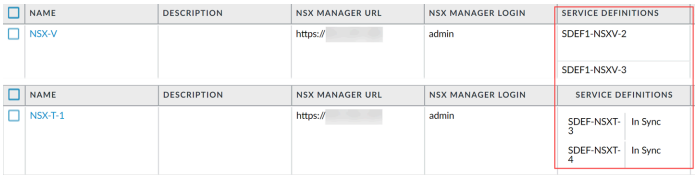

STEP 6 | Verify that your NSX-T configuration is present on NSX-T Manager.

1. Log in to NSX-T Manager.
2. Select **System** > **Service Deployments** > **Catalog**.
3. Confirm that your NSX-T service definition is listed.
4. Select **Security** > **Network Introspection Settings** > **Service Profiles**.
5. Confirm that your zones associated with your NSX-T template are listed.

STEP 7 | If you have not already done so, add a compute manager in NSX-T Manager. After you have verified that the Registration Status and Connection Status are up, continue below.

STEP 8 | Import the NSX-V configuration to NSX-T.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Uninstall the service instance from NSX-V.

This step will result in traffic disruption.

1. Log in to your vSphere client.
2. Select **Installation and Upgrade** > **Service Deployment**.
3. Select your service deployment.
4. Click **Delete**.
5. Click **Delete** to confirm.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Resolve Configuration issues on NSX-T Manager. While resolving configuration issues, you must take specific actions to migrate your VM-Series firewall configuration. In most cases, you can accept the recommendations presented by NSX-T Manager.

1. When resolving service insertion configuration, verify that you selected the correct service definition that you previously configured on Panorama for the VM-Series on NSX-T.

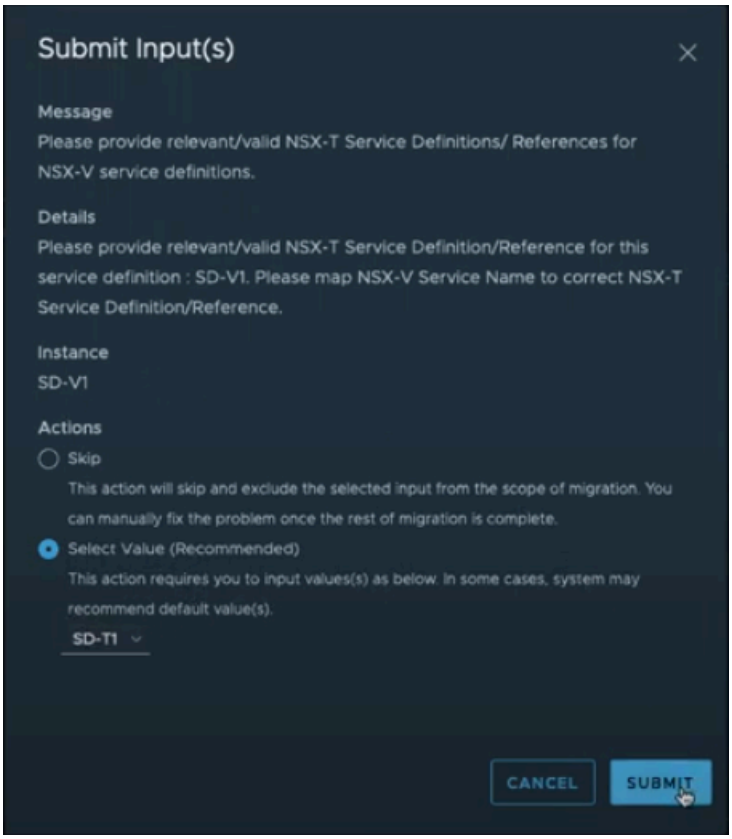

2. Continue resolving the remaining configuration.
3. Before moving to **Migrate Configuration**, you will be asked to provide a transport zone for the service insertion.
4. Map the service profiles on NSX-V to the corresponding service profiles on NSX-T.
   1. Auto_PAN_VendorTPL can be skipped.

 ©2024 Palo Alto Networks, Inc.

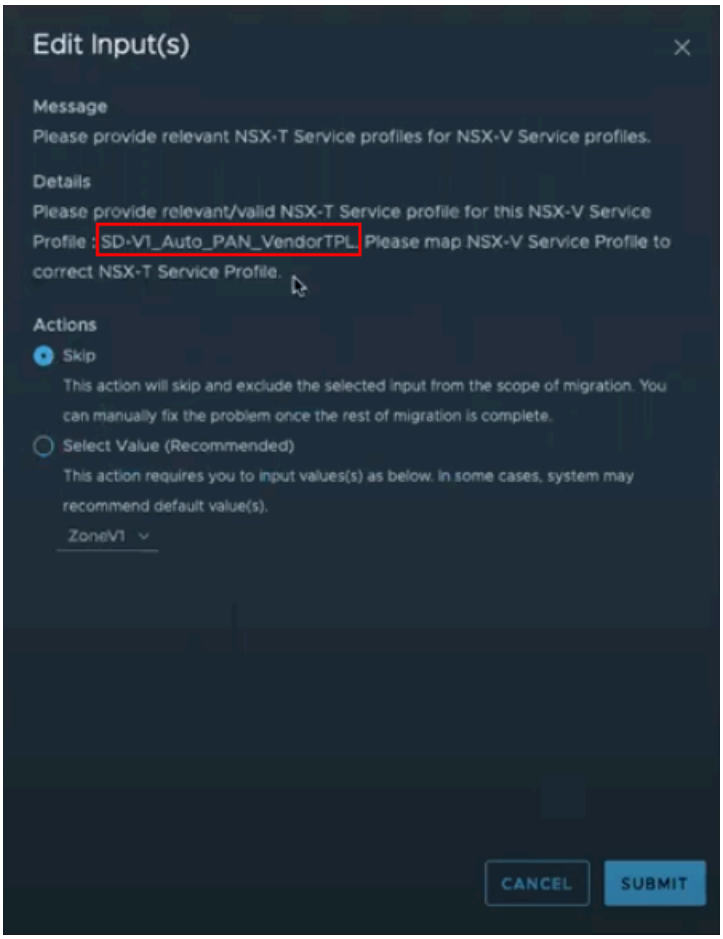

2. Map the NSX-T service profile to the corresponding NSX-V service profile.

 ©2024 Palo Alto Networks, Inc.

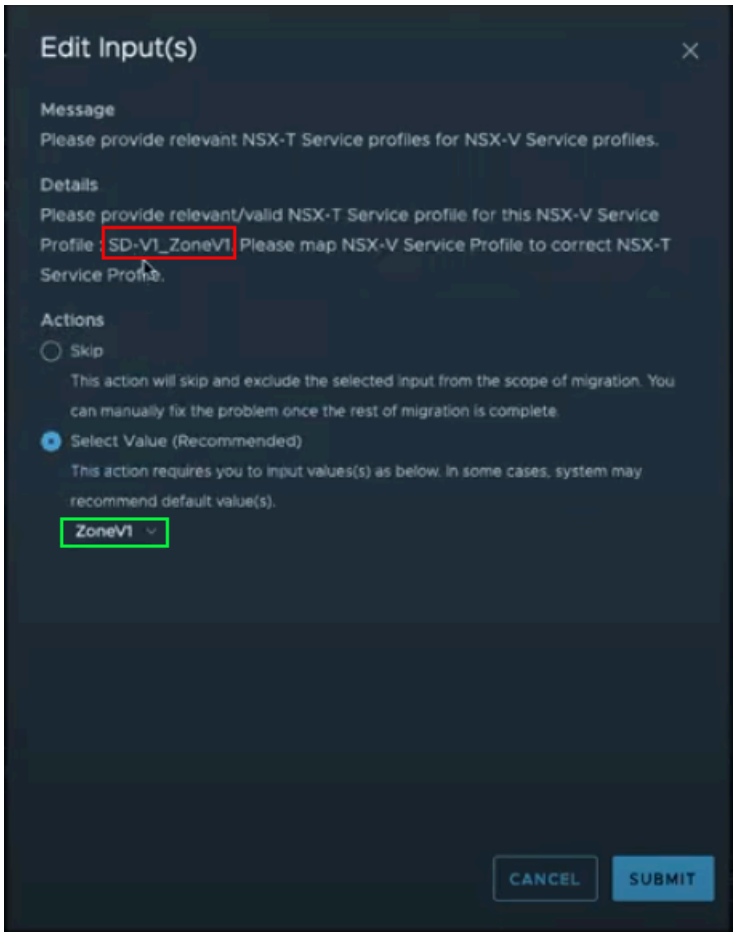

STEP 11 | Migrate the configuration.

STEP 12 | Verify that your configuration has been migrated successfully.

1. Select **Inventory** > **Groups** to verify that your IP sets and security groups are present. You can click on the security group name to see that the correct IP address are a part of the security group.
2. Select **Security** > **Network Introspection Settings** > **Service Segment** to confirm that a service segment has been created.
3. Select **Security** > **Network Introspection Settings** > **Service Chains** to confirm that a service chain has been created. Click the Profiles link in the Forward Path and Reverse Path columns to view your service profile.
4. Select **Security** > **Network Introspection (E-W)** to confirm that a traffic redirection rule has been created to direct traffic to the service profile of the VM-Series firewall.

STEP 13 | If applicable, modify and migrate edges.

STEP 14 | Configure and migrate your hosts.

 ©2024 Palo Alto Networks, Inc.

STEP 15 | Add the NSX-T tags to you existing dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click on the name of an existing NSX-V dynamic address group.
3. Click **Add Match Criteria** to display the tags from NSX-V and NSX-T.
4. Add the NSX-T tag to the dynamic address groups. If you choose not to remove the NSX-V tags, be sure to use the **OR** operator between the tags.
5. When you have added all the necessary tags, click **OK**.
6. **Commit** your changes.

STEP 16 | Launch the VM-Series Firewall on NSX-T (East-West). You do not need to create a new service segment; instead select the service segment created during migration.

Use Migration Coordinator to Move Your VM-Series from NSX-V to NSX-T

Complete the following procedure to migrate your VM-Series firewall configuration from NSX-V to NSX-T using the VMware Migration Coordinator tool in NSX-T. This tool is designed run in-place on the same hardware that is running NSX-V. By migrating your configuration, you can reuse policy and dynamic address groups already configured on Panorama. This procedure refers to information and processes published in VMware documentation as well as steps specific to PAN.

This procedure supports operations-centric NSX-V deployments only. An operations-centric deployment means that your policy rules for redirecting traffic to the VM-Series firewall were created in NSX-V Manager, not Panorama.

This procedure requires NSX-T Manager 3.1.0 or later.

It is recommended that plan for security downtime while performing this migration.

STEP 1 | Before you begin, complete the following preparatory tasks.

- Enable DRS on the cluster to be migrated.
- VMs should be deployed on shared storage and not a local datastore.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Prepare NSX-T Manager for migration.

1. Log in to the NSX-T Manager CLI via SSH.
2. Execute the following command to enable NSX Migration Coordinator.

   ```
   start service migration-coordinator
   ```

   *This will now appear in the NSX-T Manager user interface under **System** > **Migrate**.*
3. SSH in to the NSX-T Manager CLI as root or as admin.
4. Execute the following command.

   ```
   st e
   ```
5. Navigate ovf_validation.properties under /config.vmware/auth/ and make the following change.

   ```
   THIRD_PARTY_OVFS_VALIDATION_FLAG=2
   ```
6. Execute the following command.

   ```
   service cm-inventory restart
   ```

STEP 3 | Install the Panorama Plugin for VMware NSX 3.2.0 or later. See the Panorama Plugin for VMware NSX 3.2.0 Release Notes before upgrading.

STEP 4 | Enable Communication Between NSX-T Manager and Panorama.

STEP 5 | Configure an NSX-T service definition for each NSX-V service definition in your deployment. **Do not create new device groups**; instead use your existing NSX-V device groups. Using the existing device groups allows you to apply the same security policy rules used on NSX-V to the VM-Series firewalls deployed on NSX-T. If you have policy that reference a particular zone, add the same template stack from your NSX-V service definition to your NSX-T service definition. Additionally, if your device group references a particular template, ensure that you select the template stack that includes the template referenced in the device group.

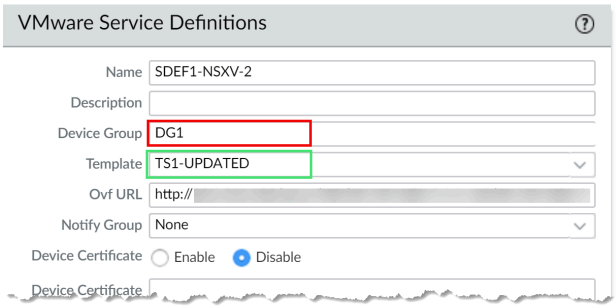

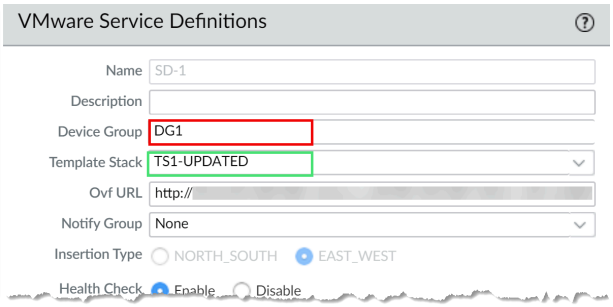

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Configure an NSX-T service manager and associate the NSX-T service definitions to the service manager.

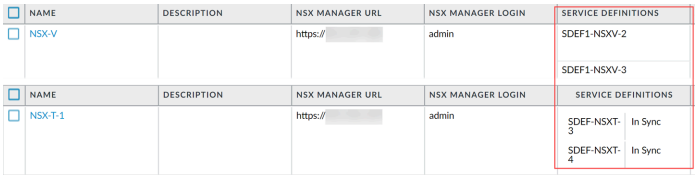

STEP 7 | Verify that your NSX-T configuration is present on NSX-T Manager.

1. Log in to NSX-T Manager.
2. Select **System** > **Service Deployments** > **Catalog**.
3. Confirm that your NSX-T service definition is listed.
4. Select **Security** > **Network Introspection Settings** > **Service Profiles**.
5. Confirm that your zones associated with your NSX-T template are listed.

Steps 8 through 11 detail the migration workflow to be executed on the migration coordinator tool on NSX-T Manager. For more information, refer to VMware migration documentation for more information.

STEP 8 | Import the NSX-V configuration to NSX-T.

STEP 9 | Uninstall the service instance from NSX-V.

This step will result in traffic disruption.

1. Log in to your vSphere client.
2. Select **Installation and Upgrade** > **Service Deployment**.
3. Select your service deployment.
4. Click **Delete**.
5. Click **Delete** to confirm.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Resolve Configuration issues on NSX-T Manager. While resolving configuration issues, you must take specific actions to migrate your VM-Series firewall configuration. In most cases, you can accept the recommendations presented by NSX-T Manager.

1. When resolving service insertion configuration, verify that you selected the correct service definition that you previously configured on Panorama for the VM-Series on NSX-T.

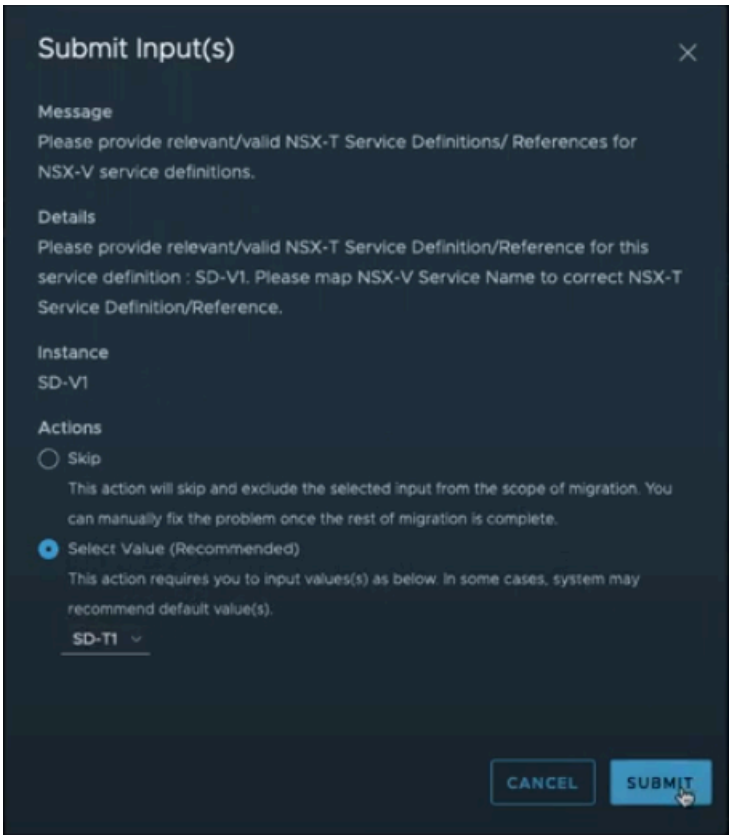

2. Continue resolving the remaining configuration.
3. Before moving to **Migrate Configuration**, you will be asked to provide a transport zone for the service insertion.
4. Map the service profiles on NSX-V to the corresponding service profiles on NSX-T.
   1. Auto_PAN_VendorTPL can be skipped.

 ©2024 Palo Alto Networks, Inc.

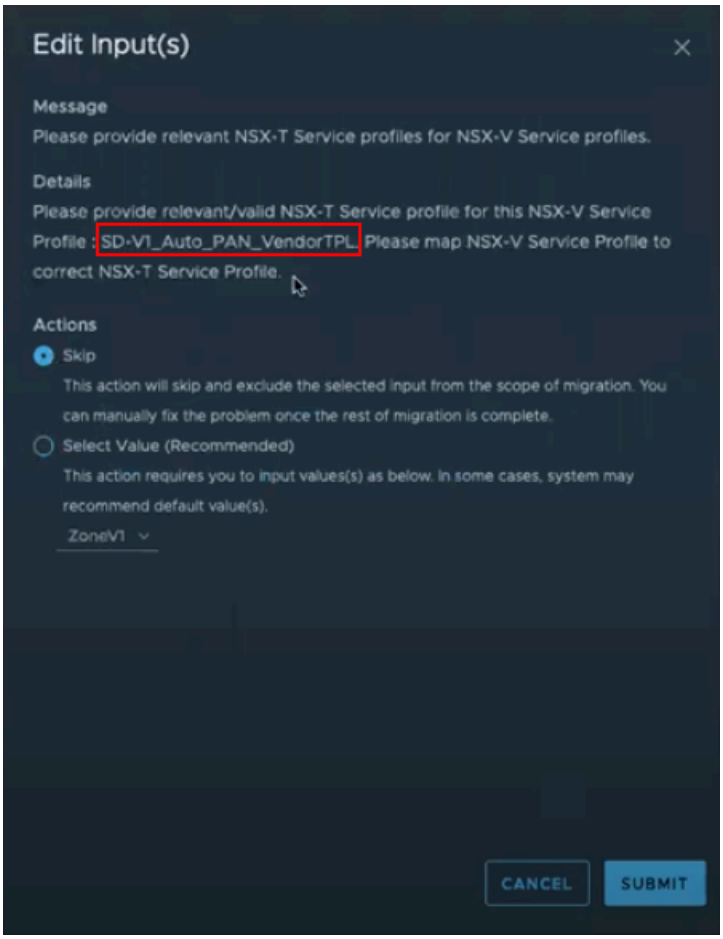

2. Map the NSX-T service profile to the corresponding NSX-V service profile.

 ©2024 Palo Alto Networks, Inc.

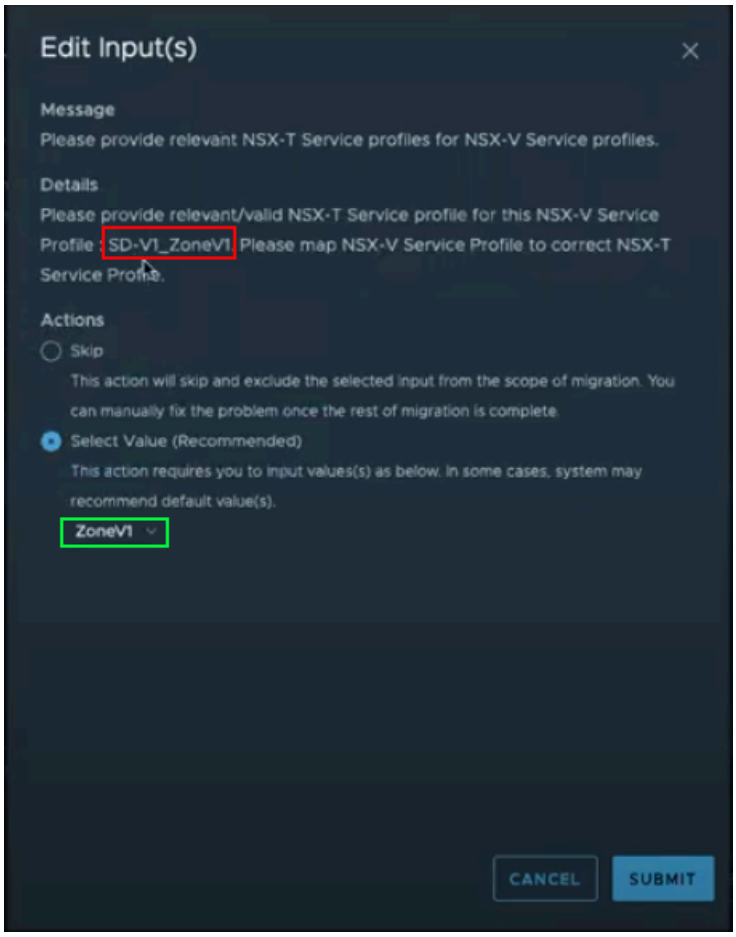

STEP 11 | Migrate the configuration.

STEP 12 | Verify that your configuration has been migrated successfully.

1. Select **Inventory** > **Groups** to verify that your IP sets and security groups are present. You can click on the security group name to see that the correct IP address are a part of the security group.
2. Select **Security** > **Network Introspection Settings** > **Service Segment** to confirm that a service segment has been created.
3. Select **Security** > **Network Introspection Settings** > **Service Chains** to confirm that a service chain has been created. Click the Profiles link in the Forward Path and Reverse Path columns to view your service profile.
4. Select **Security** > **Network Introspection (E-W)** to confirm that a traffic redirection rule has been created to direct traffic to the service profile of the VM-Series firewall.

STEP 13 | If applicable, modify and migrate edges.

STEP 14 | Configure and migrate your hosts.

 ©2024 Palo Alto Networks, Inc.

STEP 15 | Add the NSX-T tags to you existing dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click on the name of an existing NSX-V dynamic address group.
3. Click **Add Match Criteria** to display the tags from NSX-V and NSX-T.
4. Add the NSX-T tag to the dynamic address groups. If you choose not to remove the NSX-V tags, be sure to use the **OR** operator between the tags.
5. When you have added all the necessary tags, click **OK**.
6. **Commit** your changes.

STEP 16 | Launch the VM-Series Firewall on NSX-T (East-West). You do not need to create a new service segment; instead select the service segment created during migration.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on AWS

The VM-Series firewall can be deployed in the public Amazon Web Services (AWS) cloud and AWS GovCloud. It can then be configured to secure access to the applications that are deployed on EC2 instances and placed into a Virtual Private Cloud (VPC) on AWS.

- About the VM-Series Firewall on AWS
- Deployments Supported on AWS
- Deploy the VM-Series Firewall on AWS
- Panorama Orchestrated Deployments in AWS
- VM-Series Integration with an AWS Gateway Load Balancer
- High Availability for VM-Series Firewall on AWS
- Use Case: Secure the EC2 Instances in the AWS Cloud
- Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC
- Use Case: VM-Series Firewalls as GlobalProtect Gateways on AWS
- Resource Monitoring on AWS
- Auto Scale VM-Series Firewalls with the Amazon ELB Service
- List of Attributes Monitored on the AWS VPC

About the VM-Series Firewall on AWS

The Amazon Web Service (AWS) is a public cloud service that enables you to run your applications on a shared infrastructure managed by Amazon. These applications can be deployed on scalable computing capacity or EC2 instances in different AWS regions and accessed by users over the internet.

For networking consistency and ease of management of EC2 instances, Amazon offers the Virtual Private Cloud (VPC). A VPC is apportioned from the AWS public cloud, and is assigned a CIDR block from the private network space (RFC 1918). Within a VPC, you can carve public/private subnets for your needs and deploy the applications on EC2 instances within those subnets. To then enable access to the applications within the VPC, you can deploy the VM-Series firewall on an EC2 instance. The VM-Series firewall can then be configured to secure traffic to and from the EC2 instances within the VPC.

The VM-Series firewall is available in both the public AWS cloud and on AWS GovCloud. The VM-Series firewall in public AWS and AWS GovCloud supports the Bring Your Own License (BYOL) model and the hourly Pay-As-You-Go (PAYG), the usage-based licensing model that you can avail from the AWS Marketplace. For licensing details, see VM-Series Firewall Licenses for Public Clouds.

- AWS EC2 Instance Types
- VM-Series Firewall on AWS GovCloud
- VM-Series Firewall on AWS China
- VM-Series Firewall on AWS Outposts
- AWS Terminology
- Management Interface Mapping for Use with Amazon ELB
- Performance Tuning for the VM-Series Firewall on AWS

AWS EC2 Instance Types

For supported instance types, see VM-Series Models on EC2 Instances.

You can deploy the VM-Series firewall on an AWS instance size with more resources than the minimum VM-Series System Requirements. If you choose a larger instance size for the VM-Series firewall model, although the firewall only uses the max vCPU cores and memory shown in table, it does take advantage of the faster network performance that AWS provides. If you want to change the instance type on your VM-Series firewall that is licensed with the BYOL option, you must deactivate the VM before you switch the instance type to ensure that your license is valid. See Upgrade the VM-Series Model to know why.

For guidance with sizing the VM-Series firewall on AWS, refer to this article.

VM-Series Firewall on AWS GovCloud

AWS GovCloud is an isolated AWS region that meets the regulatory and compliance requirements of the US government agencies and customers.

 ©2024 Palo Alto Networks, Inc.

To secure your workloads that contain all categories of Controlled Unclassified Information (CUI) data and government-oriented, publicly available data in the AWS GovCloud (US) Region, the VM-Series firewall provides the same robust security features in the standard AWS public cloud and on AWS GovCloud. The VM-Series firewall on AWS GovCloud and the standard AWS public cloud support the same capabilities.

See AMI on AWS GovCloud to Deploy the VM-Series Firewall on AWS.

VM-Series Firewall on AWS China

The VM-Series firewall is available with the BYOL option on the AWS China Marketplace, and is available in the AWS China (Beijing) and the AWS China (Ningxia) regions. You must have an AWS China account that is separate from your global AWS account to access this image and use AWS resources on AWS China.

Make sure to review the VM-Series System Requirements before Launch the VM-Series Firewall on AWS.

VM-Series Firewall on AWS Outposts

To provide the same level of security to the workloads located on-premises as those workloads located in the AWS cloud, you can install the VM-Series firewall on AWS on an AWS Outposts rack at your on-premises location. Use the AWS marketplace BYOL AMIs for your AWS region to deploy the VM-Series firewall instances in your AWS Outposts subnets.

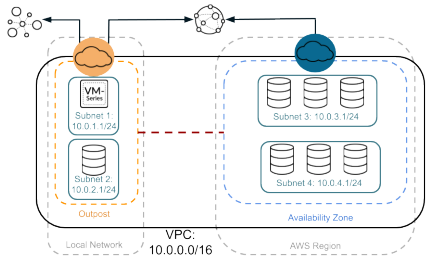

See Register the VM-Series Firewall (with auth code), to create a support account and register the VM-Series firewall on the Palo Alto Networks Customer Support website for activating your support entitlement with Palo Alto Networks.

AWS Terminology

This document assumes that you are familiar with the networking and configuration of the AWS VPC. In order to provide context for the terms used in this section, here is a brief refresher on the AWS terms (some definitions are taken directly from the AWS glossary) that are referred to in this document:

Term	Description
EC2	Elastic Compute Cloud A web service that enables you to launch and manage Linux/UNIX and Windows server instances in Amazon's data centers.

 ©2024 Palo Alto Networks, Inc.

Term	Description
AMI	Amazon Machine Image An AMI provides the information required to launch an instance, which is a virtual server in the cloud. The VM-Series AMI is an encrypted machine image that includes the operating system required to instantiate the VM-Series firewall on an EC2 instance.
ELB	Elastic Load Balancing ELB is an Amazon web service that helps you improve the availability and scalability of your applications by routing traffic across multiple Elastic Compute Cloud (EC2) instances. ELB detects unhealthy EC2 instances and reroutes traffic to healthy instances until the unhealthy instances are restored. ELB can send traffic only to the primary interface of the next hop load-balanced EC2 instance. So, to use ELB with a VM-Series firewall on AWS, the firewall must be able to use the primary interface for dataplane traffic.
ENI	Elastic Network Interface An additional network interface that can be attached to an EC2 instance. ENIs can include a primary private IP address, one or more secondary private IP addresses, a public IP address, an elastic IP address (optional), a MAC address, membership in specified security groups, a description, and a source/destination check flag.
IP address types for EC2 instances	An EC2 instance can have different types of IP addresses. • Public IP address: An IP address that can be routed across the internet. • Private IP address: A IP address in the private IP address range as defined in the RFC 1918. You can choose to manually assign an IP address or to auto assign an IP address within the range in the CIDR block for the subnet in which you launch the EC2 instance. If you are manually assigning an IP address, Amazon reserves the first four (4) IP addresses and the last one (1) IP address in every subnet for IP networking purposes. • Elastic IP address (EIP): A static IP address that you have allocated in Amazon EC2 or Amazon VPC and then attached to an instance. Elastic IP addresses are associated with your account, not with a specific instance. They are elastic because you can easily allocate, attach, detach, and free them as your needs change. An instance in a public subnet can have a Private IP address, a Public IP address, and an Elastic IP address (EIP); an instance in a private subnet will have a private IP address and optionally have an EIP.

 ©2024 Palo Alto Networks, Inc.

Term	Description
Instance type	Amazon-defined specifications that stipulate the memory, CPU, storage capacity, and hourly cost for an instance. Some instance types are designed for standard applications, whereas others are designed for CPU-intensive, memory-intensive applications, and so on.
VPC	Virtual Private Cloud An elastic network populated by infrastructure, platform, and application services that share common security and interconnection.
IGW	Internet gateway provided by Amazon. Connects a network to the internet. You can route traffic for IP addresses outside your VPC to the internet gateway.
IAM Role	Identity and Access Management Required for enabling High Availability for the VM-Series firewall on AWS. The IAM role defines the API actions and resources the application can use after assuming the role. On failover, the IAM Role allows the VM-Series firewall to securely make API requests to switch the dataplane interfaces from the active peer to the passive peer. An IAM role is also required for VM Monitoring. See List of Attributes Monitored on the AWS VPC.
Subnets	A segment of the IP address range of a VPC to which EC2 instances can be attached. EC2 instances are grouped into subnets based on your security and operational needs. There are two types of subnets: • Private subnet: The EC2 instances in this subnet cannot be reached from the internet. • Public subnet: The internet gateway is attached to the public subnet, and the EC2 instances in this subnet can be reached from the internet.
Security groups	A security group is attached to an ENI and it specifies the list of protocols, ports, and IP address ranges that are allowed to establish inbound/outbound connections on the interface. *In the AWS VPC, security groups and network ACLs control inbound and outbound traffic; security groups regulate access to the EC2 instance, while network ACLs regulate access to the subnet. Because you are deploying the VM-Series firewall, set more permissive rules in your security groups and network ACLs and allow the firewall to safely enable applications in the VPC.*

 ©2024 Palo Alto Networks, Inc.

Term	Description
Route tables	A set of routing rules that controls the traffic leaving any subnet that is associated with the route table. A subnet can be associated with only one route table.
Key pair	A set of security credentials you use to prove your identity electronically. The key pair consists of a private key and a public key. At time of launching the VM-Series firewall, you must generate a key pair or select an existing key pair for the VM-Series firewall. The private key is required to access the firewall in maintenance mode.
CloudWatch	Amazon CloudWatch is a monitoring service that allows you to collect and track metrics for the VM-Series firewalls on AWS. When enabled, the firewalls use AWS APIs to publish native PAN-OS metrics to CloudWatch.

Management Interface Mapping for Use with Amazon ELB

By default, the elastic network interface (ENI) eth0 maps to the MGT interface on the firewall and ENI eth1 maps to ethernet 1/1 on the firewall. Because the ELB can send traffic only to the primary interface of the next hop load-balanced EC2 instance, the VM-Series firewall must be able to use the primary interface for dataplane traffic.

The firewall can receive dataplane traffic on the primary interface in the following scenarios where the VM-Series firewall is behind the Amazon ELB Service (for a topology diagram, see Auto Scale VM-Series Firewalls with the Amazon ELB Service):

- The VM-Series firewall(s) is securing traffic outbound directly to the internet without the need for using a VPN link or a Direct Connect link back to the corporate network.
- The VM-Series firewall secures an internet-facing application when there is exactly one backend server, such as a web server, for each firewall. The VM-Series firewalls and web servers can scale linearly, in pairs, behind ELB.

> *At present, for use cases that require an ELB sandwich-type deployment to scale out firewalls and application layer EC2 instances, swapping the management interface will not allow you to seamlessly deploy the ELB solution. The ability to swap the management interface only partially solves the integration with ELB.*

To allow the firewall to send and receive dataplane traffic on eth0 instead of eth1, you must swap the mapping of the ENIs within the firewall such that ENI eth0 maps to ethernet 1/1 and ENI eth1 maps to the MGT interface on the firewall as shown below.

 ©2024 Palo Alto Networks, Inc.

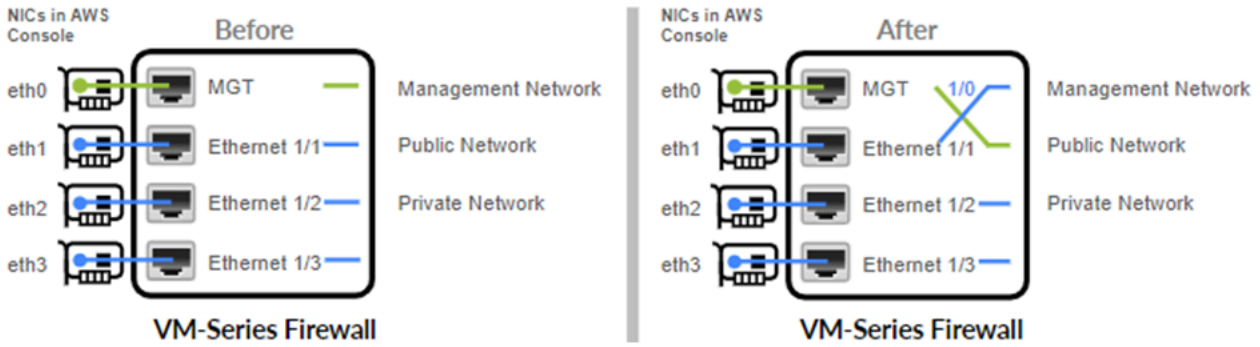

If possible, swap the management interface before you configure the firewall or define policy rules.

Swapping how the interfaces are mapped allows ELB to distribute and route traffic to healthy instances of the VM-Series firewall located in the same or different Availability Zones on AWS for increased capacity and fault tolerance.

The interface swap is only required when the VM-Series firewall is behind the Amazon ELB Service. If your requirement is to deploy the VM-Series firewalls in a traditional high availability set up, you don't need to configure the interface swap that is described in this section. Continue to High Availability for VM-Series Firewall on AWS.

To swap the interfaces, you have the following options:

- **At launch**—When you launch the firewall, you can either enter the **`mgmt-interface-swap=enable`** command in the **User data** field on the AWS management console (see Launch the VM-Series Firewall on AWS) or CLI or you can include the new **`mgmt-interface-swap`** operational command in the bootstrap configuration.
- **After launch**—After you launch the firewall, Use the VM-Series Firewall CLI to Swap the Management Interface (**`set system setting mgmt-interface-swap enable yes`** operational command) on the firewall.

  - *To prevent unpredictable behavior on the firewall, pick one method to consistently specify the interface swap setting—in the bootstrap configuration, from the CLI on the firewall, or using the Amazon EC2 **User data** field on the AWS console.*
  - *Ensure that you have access to the AWS console (management console or CLI) to view the IP address of the eth1 interface. Also, verify that the AWS Security Group rules allow connections (HTTPS and SSH) to the new management interface.*
  - *If you configured the firewall or defined policy rules before interface swap, check whether any IP address changes for eth0 or eth1 impact policy rules.*

Performance Tuning for the VM-Series Firewall on AWS

The following configurations affect performance:

- PAN-OS versions 9.0.3 and earlier support AWS C5 and M5 instance types that have Elastic Network Adapter support for SR-IOV mode by default. For more details, see Elastic Network Adapter – High Performance Network Interface for Amazon EC2.

 ©2024 Palo Alto Networks, Inc.

- PAN-OS versions 9.0.4 and later provide DPDK support for C5 and M5 instance types by default. When the firewall is in DPDK mode, it uses DPDK drivers. For a list of supported drivers, see PacketMMAP and PDK Drivers on VM-Series Firewalls and the official DPDK release notes.
- To benefit from IETF RFC 8926 (Geneve) encapsulation and improved throughput, upgrade to PAN-OS 10.0.2 or later and refer to VM-Series Integration with AWS Gateway Load Balancer.

Use the VM-Series CLI to view your DPDK settings or enable packed I/O.

View the DPDK Configuration on Your Firewall

STEP 1 | Log in to the VM-Series firewall CLI.

STEP 2 | View your DPDK configuration. If DPDK is enabled the output is as follows:

```
> show system setting dpdk-pkt-io on
```

```
Device current Packet IO mode:     DPDK
Device DPDK Packet IO capable:     yes
Device default Packet IO mode:     DPDK
```

Enable DPDK Packed I/O

STEP 1 | Log in to the VM-Series firewall CLI

STEP 2 | Enable DPDK:

```
> set system setting dpdk-pkt-io on
```

STEP 3 | Reboot the device.

On the firewall, select **Device** > **Setup** > **Operations** and select **Reboot Device**.

 ©2024 Palo Alto Networks, Inc.

Deployments Supported on AWS

The VM-Series firewall secures inbound and outbound traffic to and from EC2 instances within the AWS Virtual Private Cloud (VPC). Because the AWS VPC only supports an IP network (Layer 3 networking capabilities), the VM-Series firewall can only be deployed with Layer 3 interfaces.

- Deploy the VM-Series firewall to secure the EC2 instances hosted in the AWS Virtual Private Cloud.

  If you host your applications in the AWS cloud, deploy the VM-Series firewall to protect and safely enable applications for users who access these applications over the internet. For example, the following diagram shows the VM-Series firewall deployed in the Edge subnet to which the internet gateway is attached. The application(s) are deployed in the private subnet, which does not have direct access to the internet.

  When users need to access the applications in the private subnet, the firewall receives the request and directs it to the appropriate application, after verifying security policy and performing Destination NAT. On the return path, the firewall receives the traffic, applies security policy and uses Source NAT to deliver the content to the user. See Use Case: Secure the EC2 Instances in the AWS Cloud.

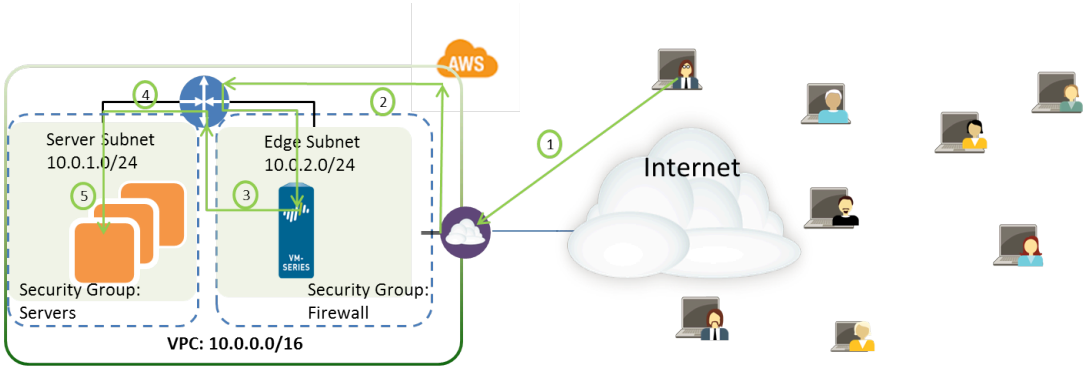

Figure 1: VM-Series for EC2 Instances

- Deploy the VM-Series firewall for VPN access between the corporate network and the EC2 instances within the AWS Virtual Private Cloud.

  To connect your corporate network with the applications deployed in the AWS Cloud, you can configure the firewall as a termination point for an IPSec VPN tunnel. This VPN tunnel allows users on your network to securely access the applications in the cloud.

  For centralized management, consistent enforcement of policy across your entire network, and for centralized logging and reporting, you can also deploy Panorama in your corporate network. If you need to set up VPN access to multiple VPCs, using Panorama allows you to group the firewalls by region and administer them with ease.

 ©2024 Palo Alto Networks, Inc.

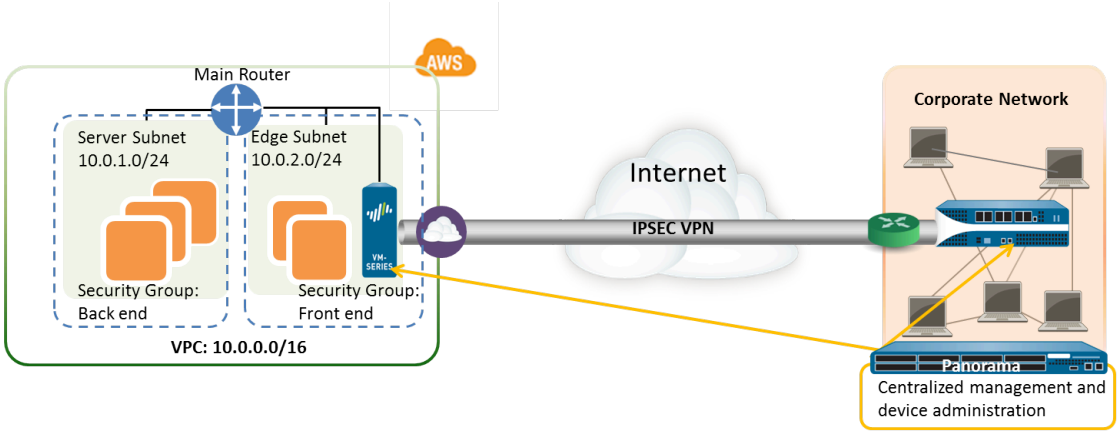

Figure 2: VM-Series for VPN Access

- Deploy the VM-Series firewall as a GlobalProtect gateway to secure access for remote users using laptops. The GlobalProtect agent on the laptop connects to the gateway, and based on the request, the gateway either sets up a VPN connection to the corporate network or routes the request to the internet. To enforce security compliance for users on mobile devices (using the GlobalProtect App), the GlobalProtect gateway is used in conjunction with the GlobalProtect Mobile Security Manager. The GlobalProtect Mobile Security Manager ensures that mobile devices are managed and configured with the device settings and account information for use with corporate applications and networks.

  In each of the use cases above, you can deploy the VM-Series firewall in an active/passive high availability (HA) pair. For information on setting up the VM-Series firewall in HA, see Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC.

- Deploy the VM-Series firewall with the Amazon Elastic Load Balancing (ELB) service, whereby the firewall can receive dataplane traffic on the primary interface in the following scenarios where the VM-Series firewall is behind the Amazon ELB:
  - The VM-Series firewall(s) is securing traffic outbound directly to the internet without the need for using a VPN link or a Direct Connect link back to the corporate network.
  - The VM-Series firewall secures an internet-facing application when there is exactly one back-end server, such as a web server, for each firewall. The VM-Series firewalls and web servers can scale linearly, in pairs, behind ELB.

  If you want to Auto Scale VM-Series Firewalls with the Amazon ELB Service, use the CloudFormation Template available in the GitHub repository repository to deploy the VM-Series in an ELB sandwich topology with an internet-facing classic ELB and an either an internal classic load balancer or an internal application load balancer (internal ELB).

 ©2024 Palo Alto Networks, Inc.

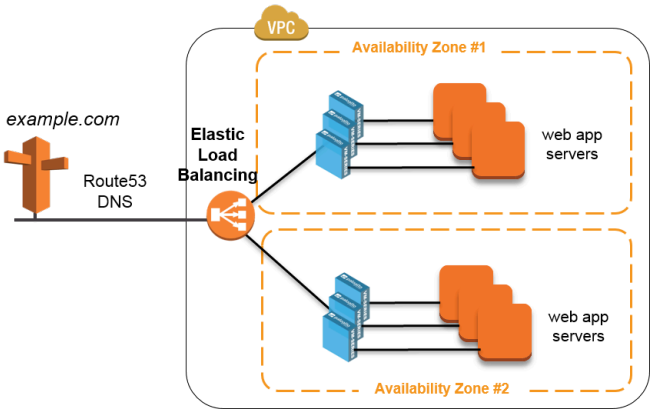

Figure 3: VM-Series with ELB

You cannot configure the firewall to send and receive dataplane traffic on eth0 when the firewall is in front of ELB. The VM-Series firewall must be placed behind the Amazon ELB.

You can either Use the VM-Series Firewall CLI to Swap the Management Interface or enable it on bootstrap. For details, see Management Interface Mapping for Use with Amazon ELB.

If you want to deploy a load balancer sandwich topology, see Auto Scale VM-Series Firewalls with the Amazon ELB Service.

In addition to the links above that are covered under the Palo Alto Networks official support policy, Palo Alto Networks provides Community supported templates in the Palo Alto Networks GitHub repository that allow you to explore the solutions available to jumpstart your journey into cloud automation and scale on AWS. See AWS Transit VPC for a hub and subscribing VPC deployment that enables you to secure traffic between VPCs, between a VPC and an on-prem/hybrid cloud resource, and secure outbound traffic to the internet.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on AWS

- Obtain the AMI
- Planning Worksheet for the VM-Series in the AWS VPC
- Launch the VM-Series Firewall on AWS
- Launch the VM-Series Firewall on AWS Outpost
- Create a Custom Amazon Machine Image (AMI)
- Encrypt EBS Volume for the VM-Series Firewall on AWS
- Use the VM-Series Firewall CLI to Swap the Management Interface
- Enable CloudWatch Monitoring on the VM-Series Firewall

Obtain the AMI

Get the Amazon Machine Image for the public AWS cloud and the AWS GovCloud from the respective Marketplace.

- AMI in the Public AWS Cloud
- AMI on AWS GovCloud
- Get the VM-Series Firewall Amazon Machine Image (AMI) ID

AMI in the Public AWS Cloud

The AMI for the VM-Series firewall is available in the AWS Marketplace for both the Bring Your Own License (BYOL) and the usage-based pricing options.

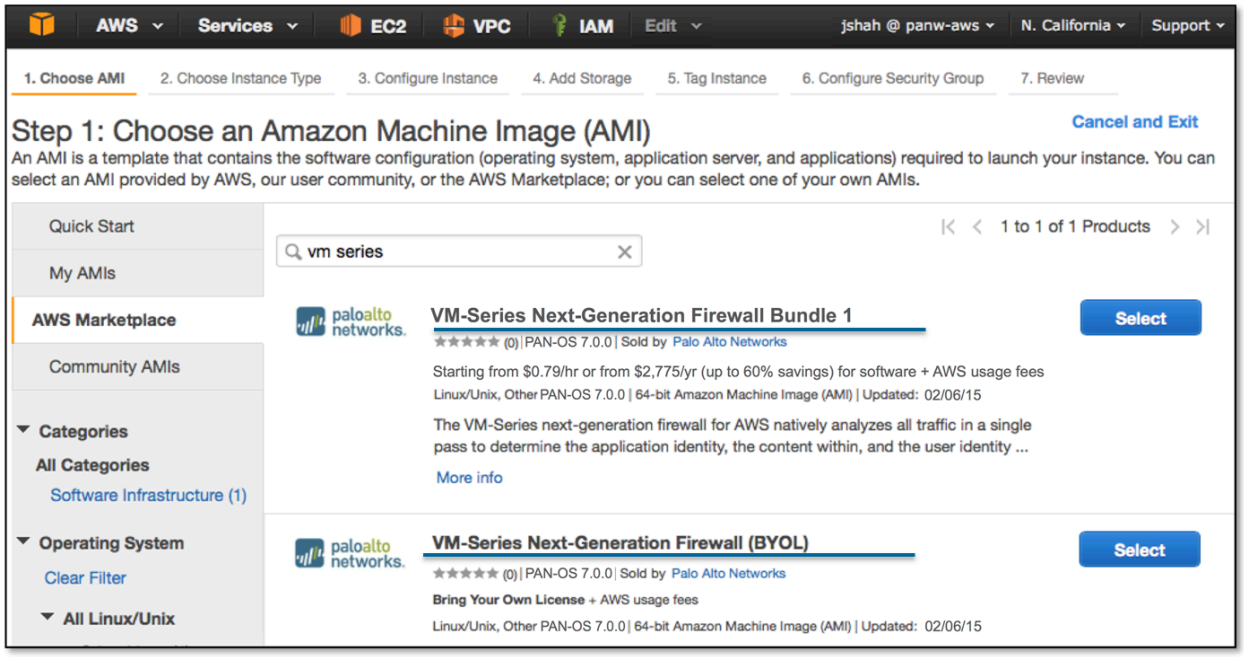

For purchasing licenses with the BYOL option, contact your Palo Alto Networks sales engineer or reseller.

 ©2024 Palo Alto Networks, Inc.

AMI on AWS GovCloud

The Bring Your Own License (BYOL) model and the usage-based model of the VM-Series firewall is available on the AWS GovCloud Marketplace.

With a GovCloud account, you can search for Palo Alto Networks and find the AMIs for the VM-Series firewall on the Marketplace. Make sure to review the supported EC2 instance types before you launch the firewall. For details, see Launch the VM-Series Firewall on AWS.

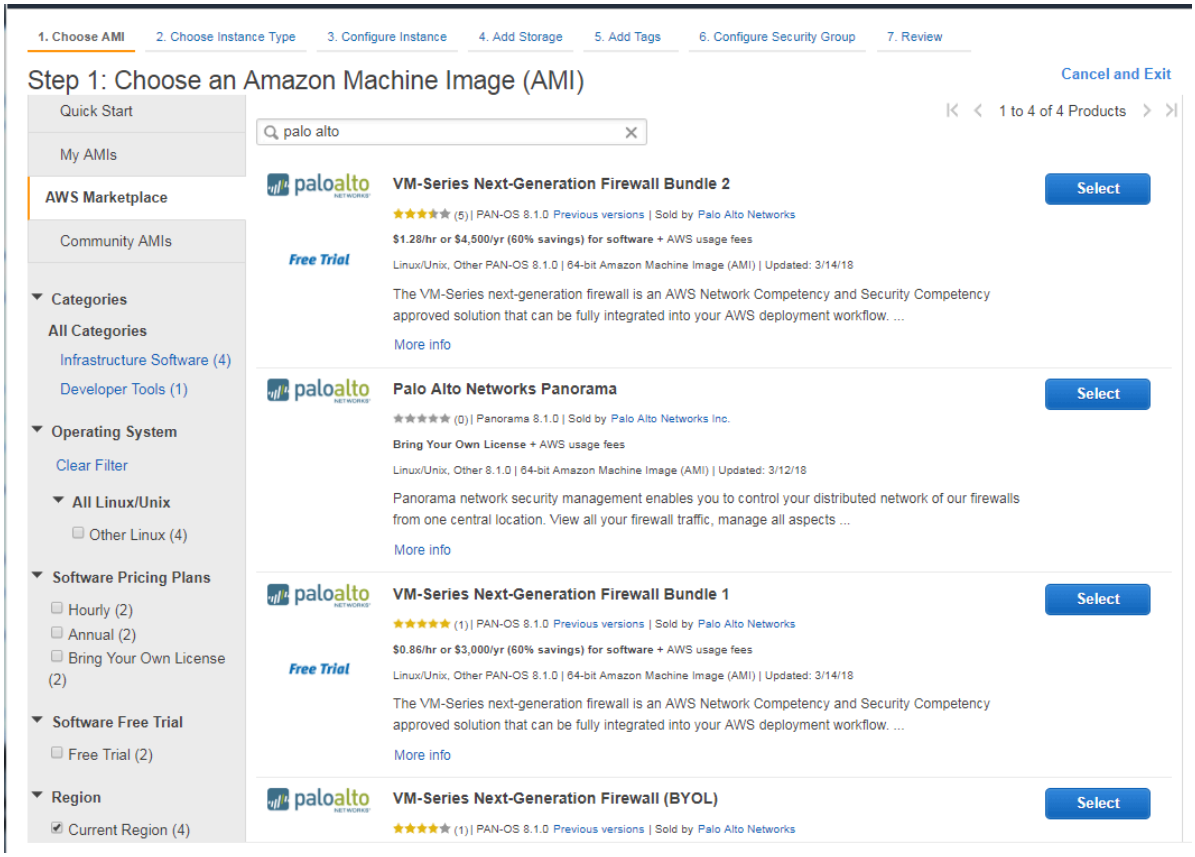

Table 1: Review System Requirements and Limitations for VM-Series on AWS

Requirement	Details
EC2 instance types	The EC2 instance type you select must meet the VM-Series System Requirements for the VM-Series firewall model. If you deploy the VM-Series firewall on an EC2 instance type that does not meet these requirements, the firewall will boot into maintenance mode *To support VM Monitoring and high availability on AWS, the VM-Series firewall must be able to directly reach the AWS API service endpoints without any proxy servers between the firewall management interface and the AWS API endpoints (such as ec2.us-west-2.amazonaws.com).*

 ©2024 Palo Alto Networks, Inc.

Requirement	Details
Amazon Elastic Block Storage (EBS)	The VM-Series firewall must use the Amazon Elastic Block Storage (EBS) volume for storage. EBS optimization provides an optimized configuration stack and additional, dedicated capacity for Amazon EBS I/O.
Networking	Because the AWS only supports Layer 3 networking capabilities, the VM-Series firewall can only be deployed with Layer 3 interfaces. Layer 2 interfaces, virtual wire, VLANs, and subinterfaces are not supported on the VM-Series firewall deployed in the AWS VPC.
Interfaces	Support for a total of eight interfaces is available—one management interface and a maximum of seven Elastic Network Interfaces (ENIs) for data traffic. The VM-Series firewall does not support hot attachment of ENIs; to detect the addition or removal of an ENI you must reboot the firewall. *Your EC2 instance type selection determines the total number of ENIs you can enable. For example, the c3.8xlarge supports eight (8) ENIs.*
Support entitlement and Licenses	For the Bring Your Own License model, a support account and a valid VM-Series license are required to obtain the Amazon Machine Image (AMI) file, which is required to install the VM-Series firewall in the AWS VPC. The licenses required for the VM-Series firewall—capacity license, support license, and subscriptions for Threat Prevention, URL Filtering, WildFire, etc—must be purchased from Palo Alto Networks. To purchase the licenses for your deployment, contact your sales representative. See VM-Series Firewall Licenses for Public Clouds. For the usage-based licensing model, hourly and annual pricing bundles can be purchased and billed directly to AWS. You must however, register your support entitlement with Palo Alto Networks. For details see, Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code).

Get the VM-Series Firewall Amazon Machine Image (AMI) ID

Use the following instructions to find the AMI ID for the VM-Series firewall that matches the PAN-OS version, license type, and AWS region in which you want to launch the VM-Series firewall.

STEP 1 | Install AWS CLI on the client that you are using to retrieve the AMI ID, and login with your AWS credentials.

Refer to the AWS documentation for instructions on installing the CLI.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Find the AMI-ID with the following CLI command.

```
aws ec2 describe-images --filters "Name=product-
code,Values=<license-type-value>" Name=name,Values=PA-VM-AWS*<PAN-
OS-version>*  --region <region> --output json
```

You need to replace the value in the angle brackets <> with the relevant information as shown below:

- Use the VM-Series product code for each license type. The values are:
  - Bundle 1—

    ```
    e9yfvyj3uag5uo5j2hjikv74n
    ```

  - Bundle 2—

    ```
    hd44w1chf26uv4p52cdynb2o
    ```

  - BYOL—

    ```
    6njl1pau431dv1qxipg63mvah
    ```

- Use the PAN-OS version— 10.0. If there are multiple feature releases within a PAN-OS version all the AMI-IDs are listed for you. For example, in 9.0.x, you will view a listing of the AMI IDs for PAN-OS versions 9.0, 9.0.3.xfr, 9.0.5.xfr, and 9.0.6, and you can use the AMI-ID for the PAN-OS version you need.
- Get the AWS region details from: https://docs.aws.amazon.com/general/latest/gr/rande.html.

For example: To find the AMI-ID for the VM-Series Bundle 1 for PAN-OS 10.0.0 in US California region, the CLI command is:

```
      aws ec2 describe-images --filters "Name=product-
code,Values=e9yfvyj3uag5uo5j2hjikv74n" "Name=name,Values=PA-VM-
AWS*10.0*"  --region us-west-1 --output json
```

The output is:

```
      {
           "ProductCodes": [
               {
                    "ProductCodeId": "e9yfvyj3uag5uo5j2hjikv74n",
```

 ©2024 Palo Alto Networks, Inc.

```
                    "ProductCodeType": "marketplace"
                }
            ],
             "VirtualizationType": "hvm",
            "Hypervisor": "xen",
            "ImageOwnerAlias": "aws-marketplace",
            "EnaSupport": true,
            "SriovNetSupport": "simple",
            "ImageId": "ami-06f7a63d7481d0ded",
            "State": "available",
            "BlockDeviceMappings": [
                {
                    "DeviceName": "/dev/xvda",
                    "Ebs": {
                        "SnapshotId": "snap-0009036179b39824b",
                        "DeleteOnTermination": false,
                        "VolumeType": "gp2",
                        "VolumeSize": 60,
                        "Encrypted": false
                    }
                }
            ],
            "Architecture": "x86_64",
            "ImageLocation": "aws-marketplace/PA-VM-AWS-10.0.0-
f1260463-68e1-4bfb-bf2e-075c2664c1d7-ami-06f7a63d7481d0ded.1",
            "RootDeviceType": "ebs",
            "OwnerId": "679593333241",
            "RootDeviceName": "/dev/xvda",
```


 ©2024 Palo Alto Networks, Inc.

```
            "CreationDate": "2020-07-20T12:45:22.000Z",

            "Public": true,

            "ImageType": "machine",

            "Name": "PA-VM-AWS-10.0.0-f1260463-68e1-4bfb-
bf2e-075c2664c1d7-ami-06f7a63d7481d0ded.1"

        }
```

You can also output to a table format. For example, to see AMI for BYOL image for PAN-OS 10.0.2:

```
aws ec2 describe-images --filters "Name=product-
code,Values=6njl1pau431dv1qxipg63mvah" "Name=name,Values=PA-VM-
AWS*10.0.2*" --region us-west-1 --output table --query "Images[*].
{Name:Name,AMI:ImageId,State:State}"
```

```
-----------------------------------------------------------------------------
                          DescribeImages
                     |+-----------------------
+-----------------------------------------------------------------------------
+-------------+|     AMI     |                    Name
                   |  State  |+-----------------------
+-----------------------------------------------------------------------------
+-------------+| ami-037b90bd9b630f594| PA-VM-
AWS-10.0.2-7064e142-2859-40a4-ab62-8b0996b842e9-
ami-07a0e94019f2a2001.4 | available |+-----------------------
+-----------------------------------------------------------------------------
+-------------+
```

Planning Worksheet for the VM-Series in the AWS VPC

For ease of deployment, plan the subnets within the VPC and the EC2 instances that you want to deploy within each subnet. Before you begin, use the following table to collate the network information required to deploy and insert the VM-Series firewall into the traffic flow in the VPC:

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

Configuration Item	Value
VPC CIDR	
Security Groups	
Subnet (public) CIDR	
Subnet (private) CIDR	

 ©2024 Palo Alto Networks, Inc.

Configuration Item	Value
Subnet (public) Route Table	
Subnet (private) Route Table	
Security Groups • Rules for Management Access to the firewall (eth0/0) • Rules for access to the dataplane interfaces of the firewall • Rules for access to the interfaces assigned to the application servers.	
VM-Series firewall behind ELB	
EC2 Instance 1 (VM-Series firewall) *An EIP is only required for the dataplane interface that is attached to the public subnet.*	Subnet: Instance type: Mgmt interface IP: Mgmt interface EIP: Dataplane interface eth1/1 • Private IP: • EIP (if required): • Security Group: Dataplane interface eth1/2 • Private IP: • EIP (if required): • Security Group:
EC2 Instance 2 (Application to be secured) Repeat these set of values for additional application(s) being deployed.	Subnet: Instance type: Mgmt interface IP: Default gateway: Dataplane interface 1 • Private IP:

 ©2024 Palo Alto Networks, Inc.

Configuration Item	Value
Requirements for HA	If you are deploying the VM-Series firewalls in a high availability (active/passive) configuration, you must ensure the following: • Create an IAM role and assign the role to the VM-Series firewall when you are deploying the instance. See IAM Roles for HA. • Deploy the HA peers in the same AWS availability zone. • The active firewall in the HA pair must have at a minimum three ENIs: two dataplane interfaces and one management interface. The passive firewall in the HA pair, must have one ENI for management, and one ENI that functions as dataplane interface; you will configure the dataplane interface as an HA2 interface. *Do not attach additional dataplane interfaces to the passive firewall in the HA pair. On failover, the dataplane interfaces from the previously active firewall are moved—detached and then attached—to the now active (previously passive) firewall.*

Launch the VM-Series Firewall on AWS

If you have not already registered the capacity authcode that you received with the order fulfillment email with your support account, see Register the VM-Series Firewall. After registering, deploy the VM-Series firewall using an AMI published in the Marketplace or Create a Custom Amazon Machine Image (AMI) in the AWS VPC as follows:

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Access the AWS Console.

Log in to the AWS console and select the EC2 Dashboard.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Set up the VPC for your network needs.

Whether you launch the VM-Series firewall in an existing VPC or you create a new VPC, the VM-Series firewall must be able to receive traffic from the EC2 instances and perform inbound and outbound communication between the VPC and the internet.

Refer to the AWS VPC documentation for instructions on creating a VPC and setting it up for access.

For an example with a complete workflow, see Use Case: Secure the EC2 Instances in the AWS Cloud.

1. Create a new VPC or use an existing VPC. Refer to the AWS Getting Started documentation.
2. Verify that the network and security components are defined suitably.
   - Enable communication to the internet. The default VPC includes an internet gateway, and if you install the VM-Series firewall in the default subnet it has access to the internet.
   - Create subnets. Subnets are segments of the IP address range assigned to the VPC in which you can launch the EC2 instances. The VM-Series firewall must belong to the public subnet so that it can be configured to access the internet.
   - Create security groups as needed to manage inbound and outbound traffic from the EC2 instances/subnets.
   - Add routes to the route table for a private subnet to ensure that traffic can be routed across subnets and security groups in the VPC, as applicable.
3. If you want to deploy a pair of VM-Series firewalls in HA, you must define IAM Roles for HA before you can Configure Active/Passive HA on AWS.
4. (Optional) If you are using bootstrapping to perform the configuration of your VM-Series firewall, refer to Bootstrap the VM-Series Firewall on AWS. For more information about bootstrapping, see Bootstrap the VM-Series Firewall and Choose a Bootstrap Method.

STEP 3 | Launch the VM-Series firewall.

Although you can add additional network interfaces (ENIs) to the VM-Series firewall when you launch, AWS releases the auto-assigned Public IP address for the management interface when you restart the firewall. Hence, to ensure connectivity to the management interface you must assign an Elastic IP address for the management interface, before attaching additional interfaces to the firewall.

If you want to conserve EIP addresses, you can assign one EIP address to the eth 1/1 interface and use this interface for both management traffic and data traffic. To restrict services

 ©2024 Palo Alto Networks, Inc.

permitted on the interface or limit IP addresses that can log in the eth 1/1 interface, attach a management profile to the interface.

1. On the EC2 Dashboard, click **Launch Instance**.
2. Select the VM-Series AMI. To get the AMI, see Obtain the AMI.
3. Launch the VM-Series firewall on an EC2 instance.
    1. Choose the **EC2 instance type** for allocating the resources required for the firewall, and click **Next**. See VM-Series System Requirements, for resource requirements.
    2. Select the VPC.
    3. Select the public subnet to which the VM-Series management interface will attach.
    4. Select **Automatically assign a public IP address**. This allows you to obtain a publicly accessible IP address for the management interface of the VM-Series firewall.

        You can later attach an Elastic IP address to the management interface; unlike the public IP address that is disassociated from the firewall when the instance is terminated, the Elastic IP address provides persistence and can be reattached to

 ©2024 Palo Alto Networks, Inc.

a new (or replacement) instance of the VM-Series firewall without the need to reconfigure the IP address wherever you might have referenced it.

5. Select **Launch as an EBS-optimized instance**.
6. Add another network interface for deployments with ELB so that you can swap the management and data interfaces on the firewall. Swapping interfaces requires a minimum of two ENIs (eth0 and eth1).
   - Expand the Network Interfaces section and click **Add Device** to add another network interface.

     Make sure that your VPC has more than one subnet so that you can add additional ENIs at launch.

     If you launch the firewall with only one ENI:
     - *The interface swap command will cause the firewall to boot into maintenance mode.*
     - *You must reboot the firewall when you add the second ENI.*
   - Expand the Advanced Details section and in the **User data** field enter **`mgmt-interface-swap=enable`** as text to perform the interface swap during launch.

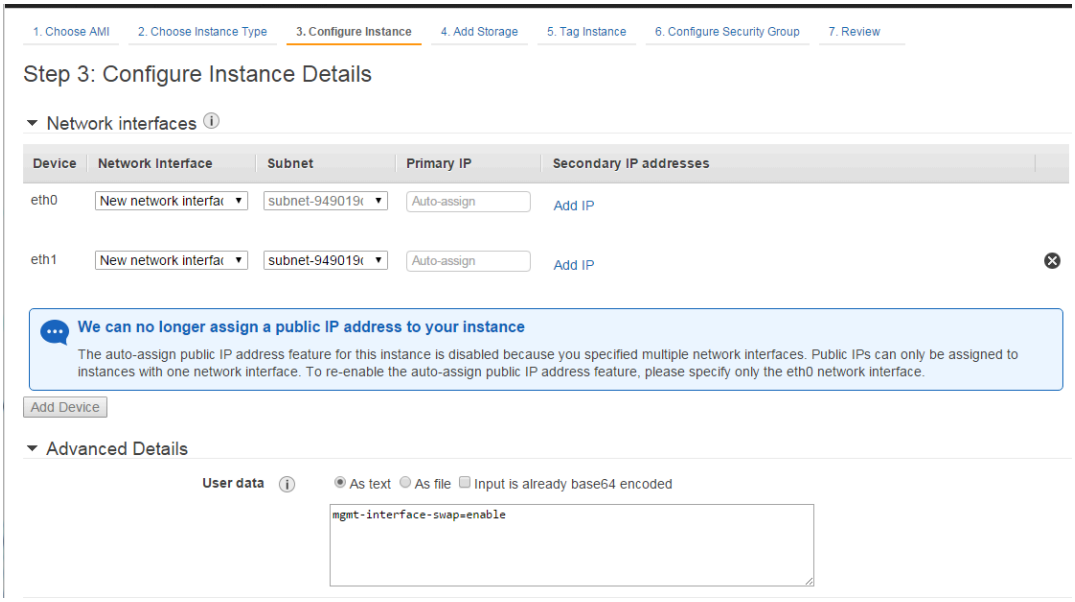

 ©2024 Palo Alto Networks, Inc.

Bootstrap Package—If you are bootstrapping the firewall with the bootstrap package, you can also enter a semicolon separator after **`mgmt-interface-swap=enable`**, then enter **`vmseries-bootstrap-aws-s3bucket=<bucketname>`**.

User Data—If you are bootstrapping with user data, enter a semicolon separator after **`mgmt-interface-swap=enable`**, and enter additional key-value pairs according to Enter a Basic Configuration as User Data (Public Clouds).

AWS Secret—If you are bootstrapping with an AWS secret, enter a semicolon separator after **`mgmt-interface-swap=enable`**, and enter the secret name as a key-value pair, as described in Step 3 of Bootstrap the VM-Series Firewall on AWS. For example:

7. Accept the default **Storage** settings. The firewall uses volume type SSD (gp2).

   This key pair is required for first time access to the firewall. It is also required to access the firewall in maintenance mode.

8. (Optional) **Tagging**. Add one or more tags to create your own metadata to identify and group the VM-Series firewall. For example, add a **Name** tag with a **Value** that helps you remember that the ENI interfaces have been swapped on this VM-Series firewall.
9. Select an existing **Security Group** or create a new one. This security group is for restricting access to the management interface of the firewall. At a minimum consider enabling https and ssh access for the management interface.
10. If prompted, select an appropriate **SSD** option for your setup.
11. Select **Review and Launch**. Review that your selections are accurate and click **Launch**.
12. Select an existing key pair or create a new one, and acknowledge the key disclaimer.
13. Download and save the private key to a safe location; the file extension is `.pem`. You cannot regenerate this key, if lost.

    It takes 5-7 minutes to launch the VM-Series firewall. You can view the progress on the EC2 Dashboard.When the process completes, the VM-Series firewall displays on the **Instances** page of the EC2 Dashboard.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Configure a new administrative password for the firewall.

On the VM-Series firewall CLI, you must configure a unique administrative password before you can access the web interface of the firewall. To log in to the CLI, you require the private key that you used to launch the firewall.

1. Use the public IP address to SSH into the Command Line Interface (CLI) of the VM-Series firewall. You will need the private key that you used or created in 3 above to access the CLI.

   If you added an additional ENI to support deployments with ELB, you must first create and assign an Elastic IP address to the ENI to access the CLI, see 6.

   If you are using PuTTY for SSH access, you must convert the .pem format to a .ppk format. See https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/putty.html
2. Enter the following command to log in to the firewall:

   ssh -i <*private_key.pem*> admin@<*public-ip_address*>
3. Configure a new password, using the following command and follow the onscreen prompts:

   configure

   set mgt-config users admin password
4. If you have a BYOL that needs to be activated, set the DNS server IP address so that the firewall can aceess the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

   set deviceconfig system dns-setting servers primary <*ip_address*>
5. Commit your changes with the command:

   commit
6. Terminate the SSH session.

STEP 5 | Shutdown the VM-Series firewall.

1. On the EC2 Dashboard, select **Instances**.
2. From the list, select the VM-Series firewall and click **Actions** > **Stop**.

STEP 6 | Create and assign an Elastic IP address (EIP) to the ENI used for management access to the firewall and reboot the VM-Series firewall.

1. Select **Elastic IPs** and click **Allocate New Address**.
2. Select **EC2-VPC** and click **Yes, Allocate**.
3. Select the newly allocated EIP and click **Associate Address**.
4. Select the **Network Interface** and the **Private IP address** associated with the management interface and click **Yes, Associate**.

STEP 7 | Create virtual network interface(s) and attach the interface(s) to the VM-Series firewall. The virtual network interfaces are called Elastic Network Interfaces (ENIs) on AWS, and serve as

 ©2024 Palo Alto Networks, Inc.

the dataplane network interfaces on the firewall. These interfaces are used for handling data traffic to/from the firewall.

You will need at least two ENIs that allow inbound and outbound traffic to/from the firewall. You can add up to seven ENIs to handle data traffic on the VM-Series firewall; check your EC2 instance type to verify the maximum number supported on it.

1. On the EC2 Dashboard, select **Network Interfaces**, and click **Create Network Interface**.
2. Enter a descriptive name for the interface.
3. Select the subnet. Use the subnet ID to make sure that you have selected the correct subnet. You can only attach an ENI to an instance in the same subnet.
4. Enter the **Private IP** address to assign to the interface or select **Auto-assign** to automatically assign an IP address within the available IP addresses in the selected subnet.
5. Select the **Security group** to control access to the dataplane network interface.
6. Click **Yes, Create**.

7. To attach the ENI to the VM-Series firewall, select the interface you just created, and click **Attach**.

8. Select the **Instance ID** of the VM-Series firewall, and click **Attach**.
9. Repeat the steps above for creating and attaching at least one more ENI to the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | (Not required for the Usage-based licensing model) Activate the licenses on the VM-Series firewall.

> *This task is not performed on the AWS management console. Access to the Palo Alto Networks support portal and the web interface of the VM-Series firewall is required for license activation.*

See Activate the License.

STEP 9 | Disable Source/Destination check on every firewall dataplane network interface(s). Disabling this option allows the interface to handle network traffic that is not destined to the IP address assigned to the network interface.

1. On the EC2 Dashboard, select the network interface, for example eth1/1, in the **Network Interfaces** tab.
2. In the **Action** drop-down, select **Change Source/Dest. Check**.

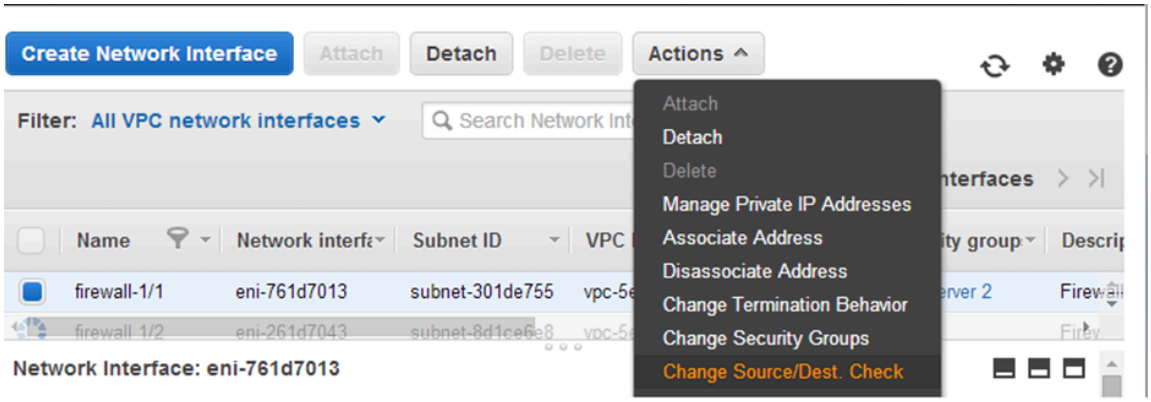

3. Click **Disabled** and **Save** your changes.
4. Repeat Steps 1-3 for each firewall dataplane interface.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

For an example configuration, see steps 14 through 17 in Use Case: Secure the EC2 Instances in the AWS Cloud.

On the application servers within the VPC, define the dataplane network interface of the firewall as the default gateway.

1. Using a secure connection (https) from your web browser, log in using the EIP address and password you assigned during initial configuration (https://<Elastic_IP address>). You will see a certificate warning; that is okay. Continue to the web page.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example VM_Series_untrust, and then click **OK**.
   - On the **IPv4** tab, select either **Static** or **DHCP Client**.

     If using the **Static** option, click **Add** in the IP section, and enter the IP address and network mask for the interface, for example 10.0.0.10/24.

     Make sure that the IP address matches the ENI IP address that you assigned earlier.

     If using DHCP, select **DHCP Client**; the private IP address that you assigned to the ENI in the AWS management console will be automatically acquired.
4. Click the link for **ethernet 1/2** and configure as follows:
   - **Interface Type**: **Layer3**
   - Security Zone: VM_Series_trust
   - **IP address**: Select the **Static** or **DHCP Client** radio button.

     For static, click **Add** in the IP section, and enter the IP address and network mask for the interface. Make sure that the IP address matches the attached ENI IP address that you assigned earlier.
5. Click **Commit**. Verify that the link state for the interfaces are up.

*For DHCP, clear the **Automatically create default route to default gateway provided by server** check box. For an interface that is attached to the private subnet in the VPC, disabling this option ensures that traffic handled by this interface does not flow directly to the internet gateway on the VPC.*

 ©2024 Palo Alto Networks, Inc.

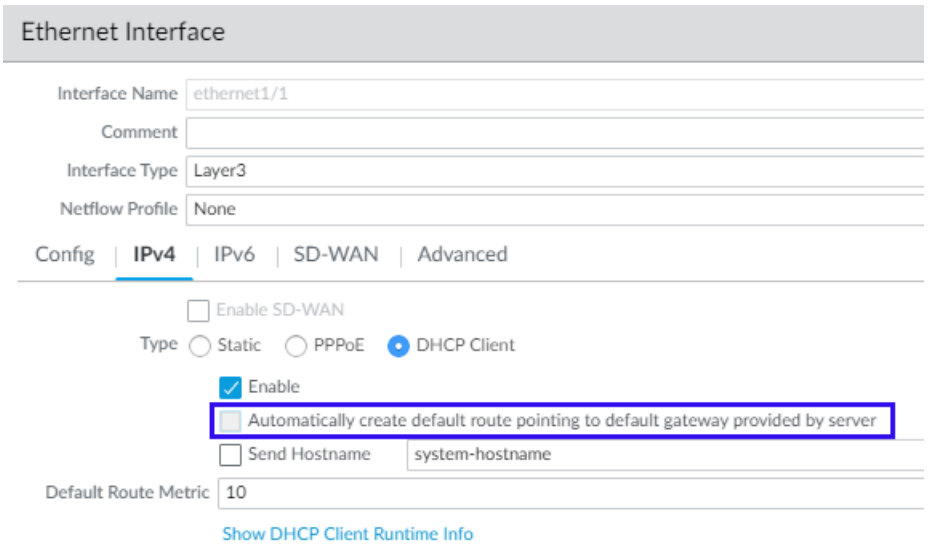

STEP 11 | Create NAT rules to allow inbound and outbound traffic from the servers deployed within the VPC.

1. Select **Policies** > **NAT** on the web interface of the firewall.
2. Create a NAT rule to allow traffic from the dataplane network interface on the firewall to the web server interface in the VPC.
3. Create a NAT rule to allow outbound access for traffic from the web server to the internet.

STEP 12 | Create security policies to allow/deny traffic to/from the servers deployed within the VPC.

1. Select **Policies** > **Security** on the web interface of the firewall.
2. Click **Add**, and specify the zones, applications and logging options that you would like to execute to restrict and audit traffic traversing through the network.

STEP 13 | Commit the changes on the firewall.

Click **Commit**.

STEP 14 | Verify that the VM-Series firewall is securing traffic and that the NAT rules are in effect.

1. Select **Monitor** > **Logs** > **Traffic** on the web interface of the firewall.
2. View the logs to make sure that the applications traversing the network match the security policies you implemented.

Launch the VM-Series Firewall on AWS Outpost

Follow this procedure to deploy the VM-Series firewall on an AWS Ouptost rack. If you have not already registered the capacity authcode that you received with the order fulfillment email, with your support account, see Register the VM-Series Firewall.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Access the AWS Outposts Console.

STEP 2 | Extend your VPC to include your AWS Outpost rack.

The VM-Series firewall must be able to receive traffic from the EC2 instances and perform inbound and outbound communication between the VPC and the internet.

Refer to the AWS Outpost documentation for instructions on connecting your Outpost to your VPC.

1. Verify that the network and security components are defined suitably.
   - Enable communication to the internet. The Outpost requires a local gateway to connect to your local LAN and the internet.
   - Create an Outpost subnet.
   - Create security groups as needed to manage inbound and outbound traffic from the EC2 instances/subnets.
   - Add routes to the route table for a private subnet to ensure that traffic can be routed across subnets and security groups in the VPC, as applicable.
2. If you want to deploy a pair of VM-Series firewalls in HA, you must define IAM Roles for HA before you can configure High Availability for VM-Series Firewall on AWS.
3. (Optional) If you are using bootstrapping to perform the configuration of your VM-Series firewall, refer to Bootstrap the VM-Series Firewall on AWS. For more information about bootstrapping, see Bootstrap the VM-Series Firewall.

STEP 3 | Launch the VM-Series firewall.

Although you can add additional network interfaces (ENIs) to the VM-Series firewall when you launch, AWS releases the auto-assigned Public IP address for the management interface when you restart the firewall. Hence, to ensure connectivity to the management interface you must assign an Elastic IP address for the management interface, before attaching additional interfaces to the firewall.

If you want to conserve EIP addresses, you can assign one EIP address to the eth 1/1 interface and use this interface for both management traffic and data traffic. To restrict services

 ©2024 Palo Alto Networks, Inc.

permitted on the interface or limit IP addresses that can log in the eth 1/1 interface, attach a management profile to the interface.

1. On the EC2 Dashboard, click **Launch Instance**.
2. Select the VM-Series AMI. To get the AMI, see Obtain the AMI.
3. Launch the VM-Series firewall on an EC2 instance.
    1. Choose the **EC2 instance type** for allocating the resources required for the firewall, and click **Next**. See VM-Series System Requirements, for resource requirements.
    2. Select the VPC.
    3. Select the public subnet on Outpost to which the VM-Series management interface will attach.
    4. Select **Automatically assign a public IP address**. This allows you to obtain a publicly accessible IP address for the management interface of the VM-Series firewall.

        You can later attach an Elastic IP address to the management interface; unlike the public IP address that is disassociated from the firewall when the instance is terminated, the Elastic IP address provides persistence and can be reattached to

 ©2024 Palo Alto Networks, Inc.

a new (or replacement) instance of the VM-Series firewall without the need to reconfigure the IP address wherever you might have referenced it.

5. Select **Launch as an EBS-optimized instance**.
6. Add another network interface for deployments with ELB so that you can swap the management and data interfaces on the firewall. Swapping interfaces requires a minimum of two ENIs (eth0 and eth1).
   - Expand the Network Interfaces section and click **Add Device** to add another network interface.

     Make sure that your VPC has more than one subnet so that you can add additional ENIs at launch.

     If you launch the firewall with only one ENI:
     - *The interface swap command will cause the firewall to boot into maintenance mode.*
     - *You must reboot the firewall when you add the second ENI.*
   - Expand the Advanced Details section and in the User data field enter **`mgmt-interface-swap=enable`** as text to perform the interface swap during launch.

     *If you are bootstrapping the firewall, you can also enter **`vmseries-bootstrap-aws-s3bucket=<bucketname>`** with a comma separator after **`mgmt-interface-swap=enable`**.*

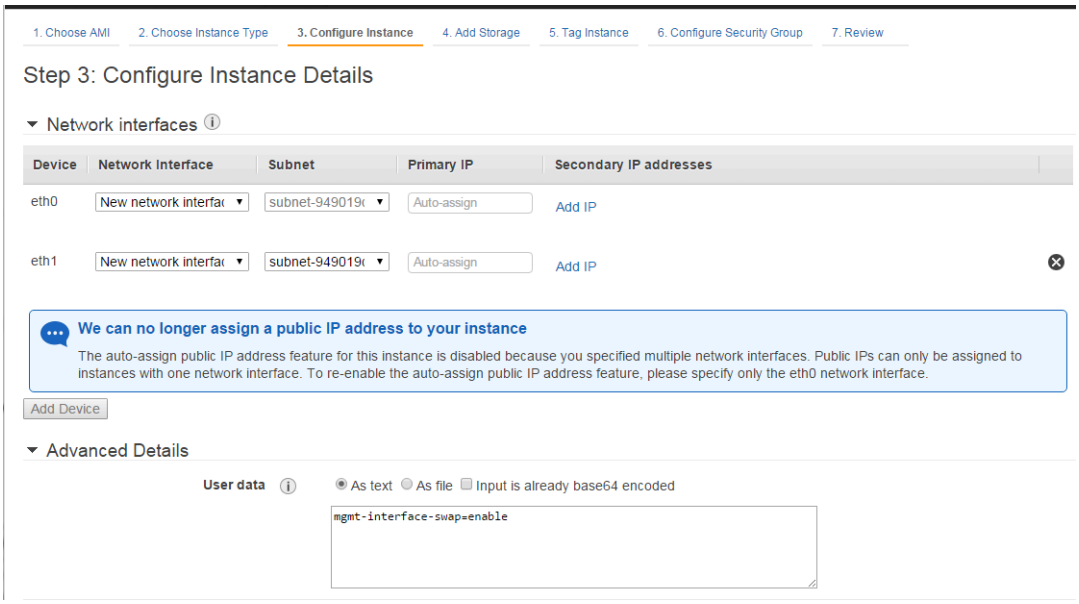

7. Accept the default **Storage** settings. The firewall uses volume type SSD (gp2)

   This key pair is required for first time access to the firewall. It is also required to access the firewall in maintenance mode.
8. (Optional) **Tagging**. Add one or more tags to create your own metadata to identify and group the VM-Series firewall. For example, add a **Name** tag with a **Value** that

 ©2024 Palo Alto Networks, Inc.

helps you remember that the ENI interfaces have been swapped on this VM-Series firewall.

9. Select an existing **Security Group** or create a new one. This security group is for restricting access to the management interface of the firewall. At a minimum consider enabling https and ssh access for the management interface.
10. If prompted, select an appropriate **SSD** option for your setup.
11. Select **Review and Launch**. Review that your selections are accurate and click **Launch**.
12. Select an existing key pair or create a new one, and acknowledge the key disclaimer.
13. Download and save the private key to a safe location; the file extension is `.pem`. You cannot regenerate this key, if lost.

    It takes 5-7 minutes to launch the VM-Series firewall. You can view the progress on the EC2 Dashboard.When the process completes, the VM-Series firewall displays on the **Instances** page of the EC2 Dashboard.

STEP 4 | Configure a new administrative password for the firewall.

On the VM-Series firewall CLI, you must configure a unique administrative password before you can access the web interface of the firewall. To log in to the CLI, you require the private key that you used to launch the firewall.

1. Use the public IP address to SSH into the Command Line Interface (CLI) of the VM-Series firewall. You will need the private key that you used or created in 3 above to access the CLI.

   If you added an additional ENI to support deployments with ELB, you must first create and assign an Elastic IP address to the ENI to access the CLI, see 6.

   If you are using PuTTY for SSH access, you must convert the .pem format to a .ppk format. See https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/putty.html
2. Enter the following command to log in to the firewall:

   ssh -i <*private_key.pem*> admin@<*public-ip_address*>
3. Configure a new password, using the following command and follow the onscreen prompts:

   configure

   set mgt-config users admin password
4. If you have a BYOL that needs to be activated, set the DNS server IP address so that the firewall can aceess the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

   set deviceconfig system dns-setting servers primary <*ip_address*>
5. Commit your changes with the command:

   commit
6. Terminate the SSH session.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Shutdown the VM-Series firewall.

1. On the EC2 Dashboard, select **Instances**.
2. From the list, select the VM-Series firewall and click **Actions** > **Stop**.

STEP 6 | Create and assign an Elastic IP address (EIP) to the ENI used for management access to the firewall and reboot the VM-Series firewall.

1. Select **Elastic IPs** and click **Allocate New Address**.
2. Select **EC2-VPC** and click **Yes, Allocate**.
3. Select the newly allocated EIP and click **Associate Address**.
4. Select the **Network Interface** and the **Private IP address** associated with the management interface and click **Yes, Associate**.

STEP 7 | Create virtual network interface(s) and attach the interface(s) to the VM-Series firewall. The virtual network interfaces are called Elastic Network Interfaces (ENIs) on AWS, and serve as the dataplane network interfaces on the firewall. These interfaces are used for handling data traffic to/from the firewall.

You will need at least two ENIs that allow inbound and outbound traffic to/from the firewall. You can add up to seven ENIs to handle data traffic on the VM-Series firewall; check your EC2 instance type to verify the maximum number supported on it.

1. On the EC2 Dashboard, select **Network Interfaces**, and click **Create Network Interface**.
2. Enter a descriptive name for the interface.
3. Select the subnet. Use the subnet ID to make sure that you have selected the correct subnet. You can only attach an ENI to an instance in the same subnet.
4. Enter the **Private IP** address to assign to the interface or select **Auto-assign** to automatically assign an IP address within the available IP addresses in the selected subnet.
5. Select the **Security group** to control access to the dataplane network interface.
6. Click **Yes, Create**.

7. To attach the ENI to the VM-Series firewall, select the interface you just created, and click **Attach**.

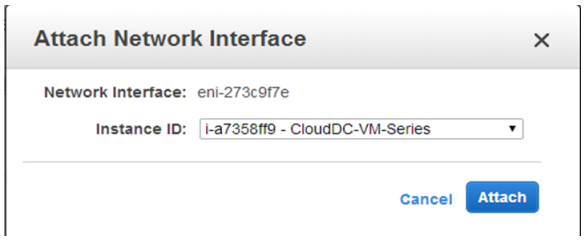

8. Select the **Instance ID** of the VM-Series firewall, and click **Attach**.
9. Repeat the steps above for creating and attaching at least one more ENI to the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | (Not required for the Usage-based licensing model) Activate the licenses on the VM-Series firewall.

This task is not performed on the AWS management console. Access to the Palo Alto Networks support portal and the web interface of the VM-Series firewall is required for license activation.

See Activate the License.

STEP 9 | Disable Source/Destination check on every firewall dataplane network interface(s). Disabling this option allows the interface to handle network traffic that is not destined to the IP address assigned to the network interface.

1. On the EC2 Dashboard, select the network interface, for example eth1/1, in the **Network Interfaces** tab.
2. In the **Action** drop-down, select **Change Source/Dest. Check**.

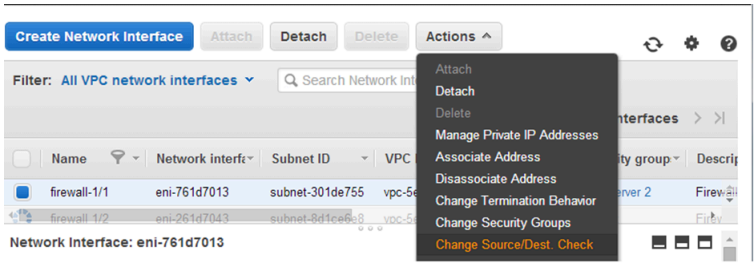

3. Click **Disabled** and **Save** your changes.
4. Repeat Steps 1-3 for each firewall dataplane interface.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

For an example configuration, see steps 14 through 17 in Use Case: Secure the EC2 Instances in the AWS Cloud.

On the application servers within the VPC, define the dataplane network interface of the firewall as the default gateway.

1. Using a secure connection (https) from your web browser, log in using the EIP address and password you assigned during initial configuration (https://<Elastic_IP address>). You will see a certificate warning; that is okay. Continue to the web page.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example VM_Series_untrust, and then click **OK**.
   - On the **IPv4** tab, select either **Static** or **DHCP Client**.

     If using the **Static** option, click **Add** in the IP section, and enter the IP address and network mask for the interface, for example 10.0.0.10/24.

     Make sure that the IP address matches the ENI IP address that you assigned earlier.

     If using DHCP, select **DHCP Client**; the private IP address that you assigned to the ENI in the AWS management console will be automatically acquired.
4. Click the link for **ethernet 1/2** and configure as follows:
   - **Interface Type**: **Layer3**
   - Security Zone: VM_Series_trust
   - **IP address**: Select the **Static** or **DHCP Client** radio button.

     For static, click **Add** in the IP section, and enter the IP address and network mask for the interface. Make sure that the IP address matches the attached ENI IP address that you assigned earlier.
5. Click **Commit**. Verify that the link state for the interfaces are up.

*For DHCP, clear the **Automatically create default route to default gateway provided by server** check box. For an interface that is attached to the private subnet in the VPC, disabling this option ensures that traffic handled by this interface does not flow directly to the internet gateway on the VPC.*

 ©2024 Palo Alto Networks, Inc.

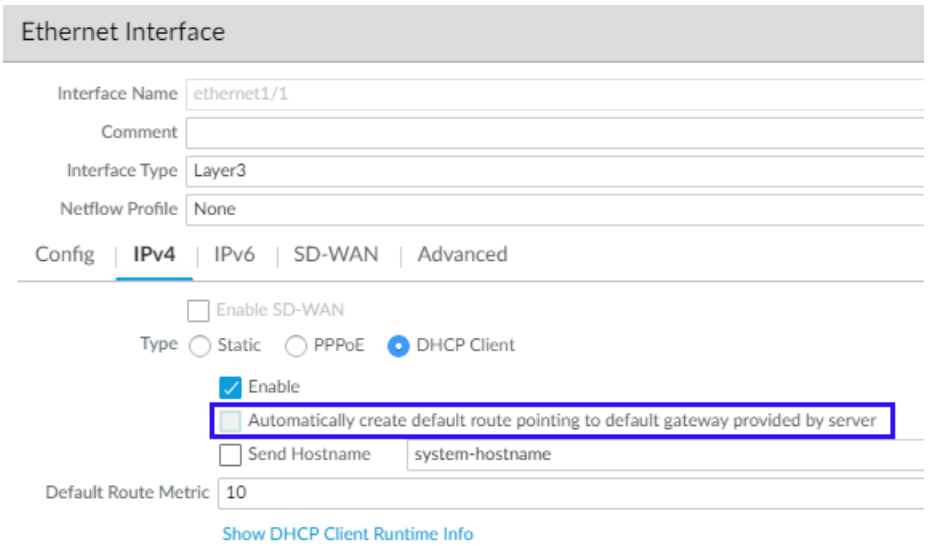

STEP 11 | Create NAT rules to allow inbound and outbound traffic from the servers deployed within the VPC.

1. Select **Policies** > **NAT** on the web interface of the firewall.
2. Create a NAT rule to allow traffic from the dataplane network interface on the firewall to the web server interface in the VPC.
3. Create a NAT rule to allow outbound access for traffic from the web server to the internet.

STEP 12 | Create security policies to allow/deny traffic to/from the servers deployed within the VPC.

1. Select **Policies** > **Security** on the web interface of the firewall.
2. Click **Add**, and specify the zones, applications and logging options that you would like to execute to restrict and audit traffic traversing through the network.

STEP 13 | Commit the changes on the firewall.

Click **Commit**.

STEP 14 | Verify that the VM-Series firewall is securing traffic and that the NAT rules are in effect.

1. Select **Monitor** > **Logs** > **Traffic** on the web interface of the firewall.
2. View the logs to make sure that the applications traversing the network match the security policies you implemented.

Create a Custom Amazon Machine Image (AMI)

A custom VM-Series AMI gives you the consistency and flexibility to deploy a VM-Series firewall with the PAN-OS version you want to use on your network instead of being restricted to using only an AMI that is published to the AWS public Marketplace or to the AWS GovCloud Marketplace. Using a custom AMI speeds up the process of deploying a firewall with the PAN-OS version of your choice because it reduces the time to provision the firewall with an AMI published on the AWS public or AWS GovCloud marketplace, and then performing software upgrades to get to the PAN-OS version you have qualified or want to use on your network. Additionally, you can

 ©2024 Palo Alto Networks, Inc.

then use the custom AMI in the Auto Scaling VM-Series Firewalls CloudFormation Templates or any other templates that you have created.

You can create a custom AMI with the BYOL, Bundle 1, or Bundle 2 licenses. The process of creating a custom AMI requires you to remove all configuration from the firewall and reset it to factory defaults, so in this workflow you'll launch a new instance of the firewall from the AWS Marketplace instead of using an existing firewall that you have fully configured.

When creating a custom AMI with a BYOL version of the firewall, you must first activate the license on the firewall so that you can access and download PAN-OS software updates to upgrade your firewall, and then deactivate the license on the firewall before you reset the firewall to factory defaults and create the custom AMI. If you do not deactivate the license, you lose the license that you applied on this firewall instance.

STEP 1 | Launch the VM-Series firewall from the Marketplace.

See 3

STEP 2 | (Only for BYOL) Activate the license.

STEP 3 | Install software updates and upgrade the firewall to the PAN-OS version you plan to use.

STEP 4 | (Only for BYOL) Deactivate the license.

STEP 5 | Perform a private data reset.

A private data reset removes all logs and restores the default configuration.

The system disks are not erased, so the content updates from Step 4 are intact.

1. Access the firewall CLI.
2. Remove all logs and restore the default configuration.

```
request system private-data-reset
```

Enter **y** to confirm.

The firewall reboots to initialize the default configuration.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Create the custom AMI.

1. Log in to the AWS Console and select the EC2 Dashboard.
2. **Stop** the VM-Series firewall.
3. Select the VM-Series firewall instance, and click **Image** > **Create Image**.

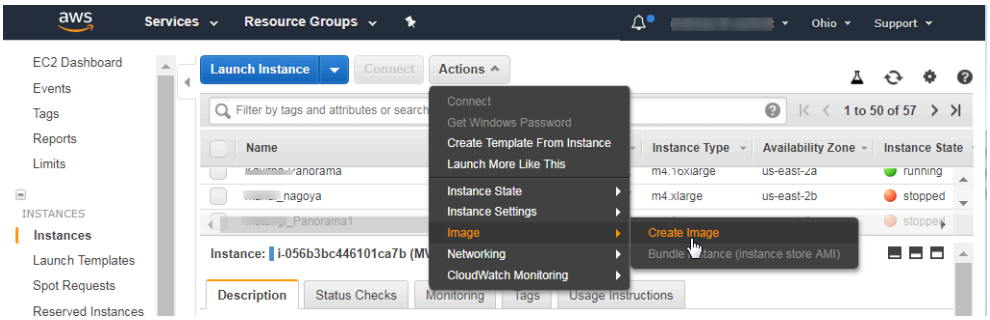

4. Enter a custom image name, and click **Create Image**.

   The disk space of 60GB is the minimum requirement.

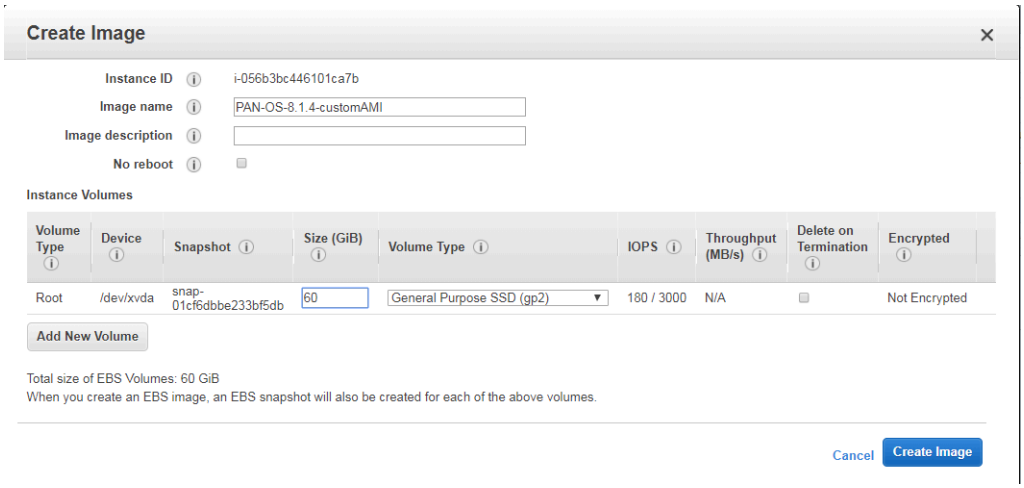

5. Verify that the custom AMI is created and has the correct product code.
   1. On the EC2 Dashboard, select **AMI**.
   2. Select the AMI that you just created. Depending on whether you selected an AMI with the BYOL, Bundle 1, or Bundle 2 licensing options, you should see one of the following **Product Codes** in the details:
      - BYOL—6njl1pau431dv1qxipg63mvah
      - Bundle 1—6kxdw3bbmdeda3o6i1ggqt4km
      - Bundle 2—806j2of0qy5osgjjixq9gqc6g

 ©2024 Palo Alto Networks, Inc.

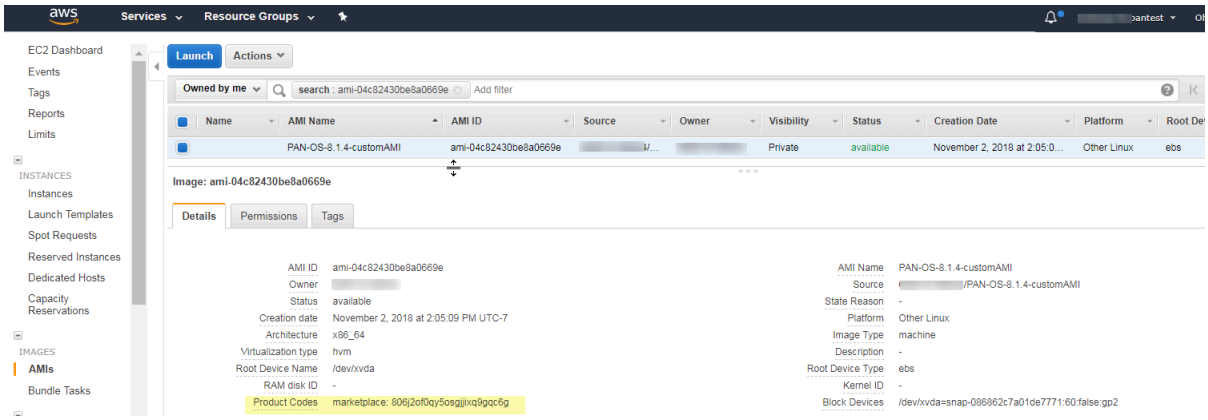

STEP 7 | Encrypt EBS Volume for the VM-Series Firewall on AWS.

If you plan to use the custom AMI with EBS encryption for an Auto Scaling VM-Series Firewalls with the Amazon ELB Service deployment, you must use the default master key for your AWS account.

STEP 8 | Configure the administrative password on the firewall.

See 4

Encrypt EBS Volume for the VM-Series Firewall on AWS

EBS encryption is available for all AWS EC2 Instance Types on which you can deploy the VM-Series firewall. To securely store data on the VM-Series firewall on AWS, you must first create a copy of an AMI that is published on the AWS public or GovCloud Marketplace, or use a custom AMI, and then encrypt the EBS volume with a customer master key (CMK) on the AWS Key Management Service (KMS). You can use the default master key for your AWS account or any CMK that you have previously created using the AWS Key Management Service, and EBS the KMS interact to ensure data security.

STEP 1 | Create an encryption key on AWS or skip this step if you want to use the default master key for your account.

You will use this key to encrypt the EBS volume on the firewall. Note that the key is region specific.

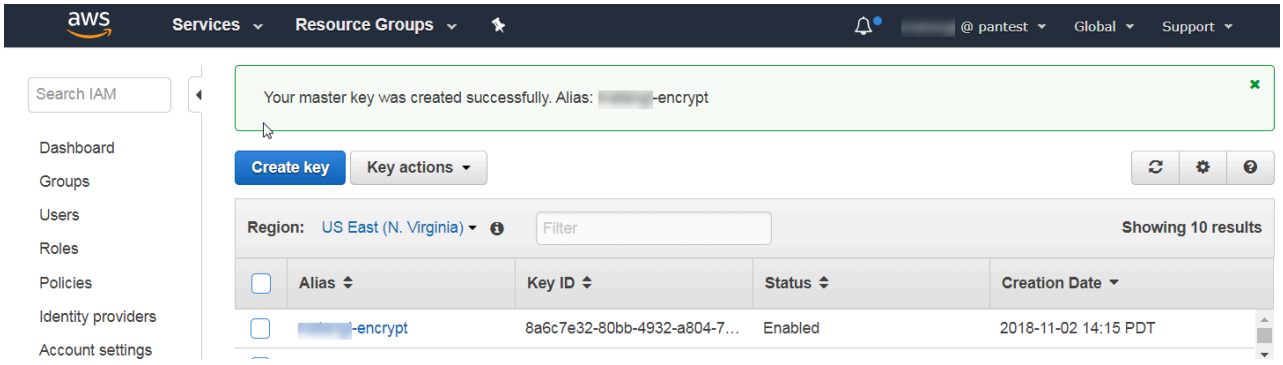

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Use the key to encrypt the EBS volume on the firewall.

You must create a copy of the AMI that you want to encrypt. You can copy an AMI that is published on the AWS public or GovCloud Marketplace, or use a custom AMI (Create a Custom Amazon Machine Image (AMI)).

1. On the EC2 Dashboard, select the AMI and **Copy AMI**.

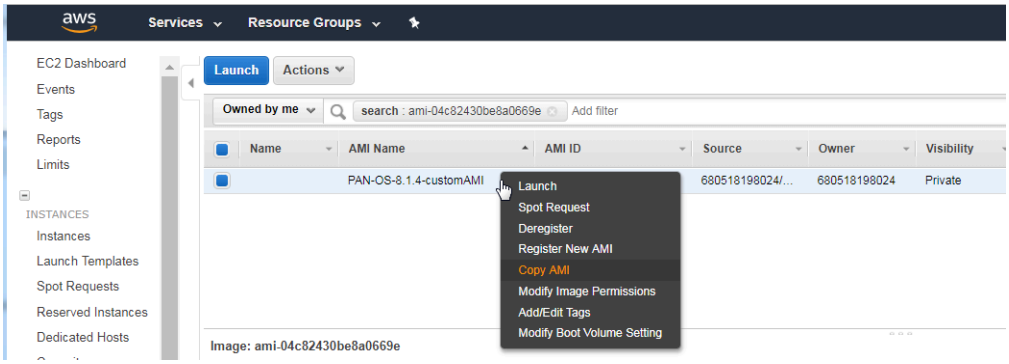

2. Set the details for the AMI.

   Make sure to select **Encrypt target EBS snapshots**.

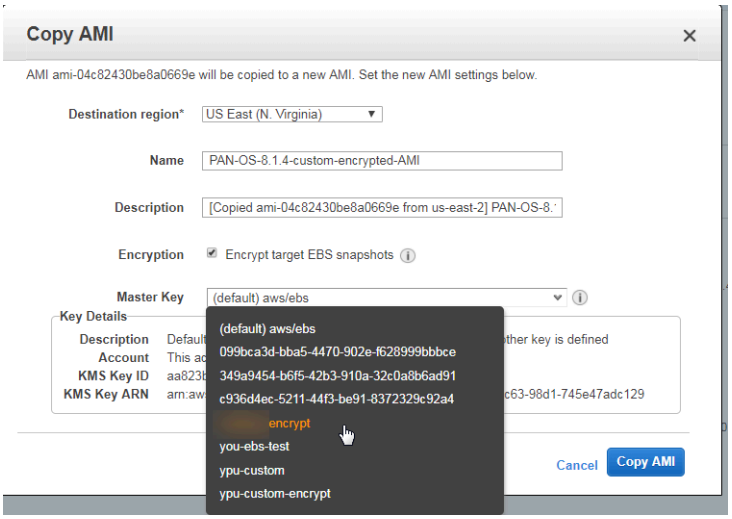

3. Select the encryption key and **Copy AMI** to create an encrypted EBS snapshot.

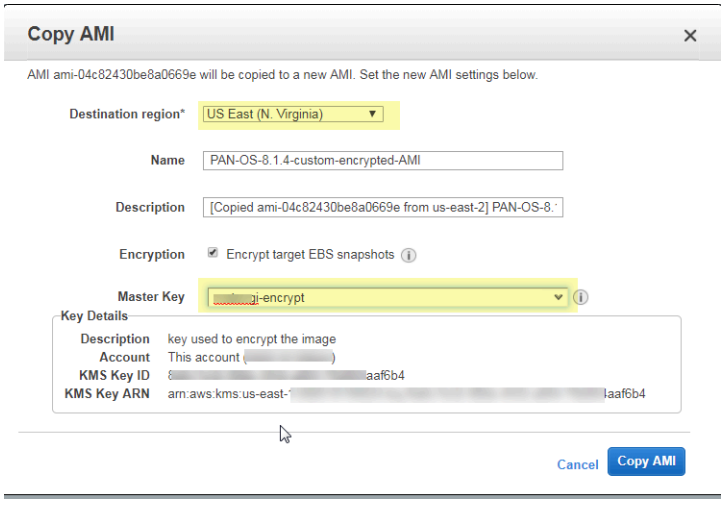

 ©2024 Palo Alto Networks, Inc.

4. Select **EC2 Dashboard** > **Snapshots** to verify that the EBS snapshot is encrypted with the key you selected above.

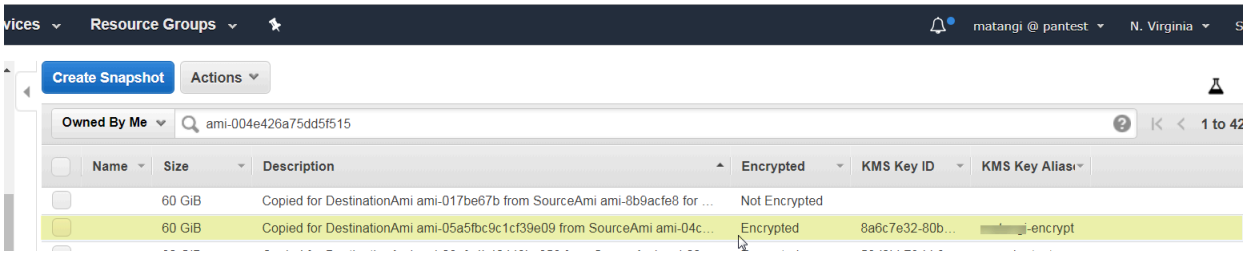

Use the VM-Series Firewall CLI to Swap the Management Interface

If you did not swap the management interface (MGT) with the dataplane interface (ethernet 1/1) when deploying the firewall, you can use the CLI to enable the firewall to receive dataplane traffic on the primary interface after launching the firewall.

STEP 1 | Complete Steps 1 through 7 in Launch the VM-Series Firewall on AWS.

Before you proceed, verify that the firewall has a minimum of two ENIs (eth0 and eth1). If you launch the firewall with only one ENI, the interface swap command will cause the firewall to boot into maintenance mode.

STEP 2 | On the EC2 Dashboard, view the IP address of the eth1 interface and verify that the AWS Security Group rules allow connections (HTTPS and SSH) to the new management interface (eth1).

STEP 3 | Log in to the VM-Series firewall CLI and enter the following command:

```
set system setting mgmt-interface-swap enable yes
```

STEP 4 | Confirm that you want to swap the interface and use the eth1 dataplane interface as the management interface.

STEP 5 | Reboot the firewall for the swap to take effect. Use the following command:

```
request restart system
```

STEP 6 | Verify that the interfaces have been swapped. Use the following command:

```
debug show vm-series interfaces all
Phoenix_interface   Base-OS_port   Base-OS_MAC        PCI-ID
       Driver
mgt(interface-swap) eth0   0e:53:96:91:ef:29   0000:00:04.0
  ixgbevf
Ethernet1/1           eth1    0e:4d:84:5f:7f:4d       0000:00:03.0
  ixgbevf
```

 ©2024 Palo Alto Networks, Inc.

Enable CloudWatch Monitoring on the VM-Series Firewall

The VM-Series firewall on AWS can publish native PAN-OS metrics to AWS CloudWatch, which you can use to monitor the firewalls. These metrics allow you to assess performance and usage patterns that you can use to take action for launching or terminating instances of the VM-Series firewalls.

The firewalls use AWS APIs to publish the metric to a *namespace*, which is the location on AWS where the metrics are collected at a specified time interval. When you configure the firewalls to publish metrics to AWS CloudWatch, there are two namespaces where you can view metrics—the primary namespace collects and aggregates the selected metric for all instances configured to use the namespace, and the secondary namespace that is automatically created with the suffix `_dimensions` allows you to filter the metrics using the hostname and AWS instance ID metadata (or *dimensions*) and get visibility into the usage and performance of individual VM-Series firewalls.

You can monitor the metric in CloudWatch or create auto scaling policies to trigger alarms and take an action to manually deploy a new instance of the firewall when the monitored metric reaches a threshold value. Refer to the AWS CloudWatch and Auto Scaling Groups (ASG) documentation on best practices for setting the alarm conditions for a scale out or scale in action.

For a description on the PAN-OS metrics that you can publish to CloudWatch, see Custom PAN-OS Metrics Published for Monitoring.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Assign the appropriate permissions for the AWS Identity and Access Management (IAM) user role that you use to deploy the VM-Series firewall on AWS.

Whether you launch a new instance of the VM-Series firewall or upgrade an existing VM-Series firewall on AWS, the IAM role associated with your instance, must have permissions to publish metrics to CloudWatch.

1. On the AWS console, select IAM.
2. Edit the IAM role to grant the following permissions:

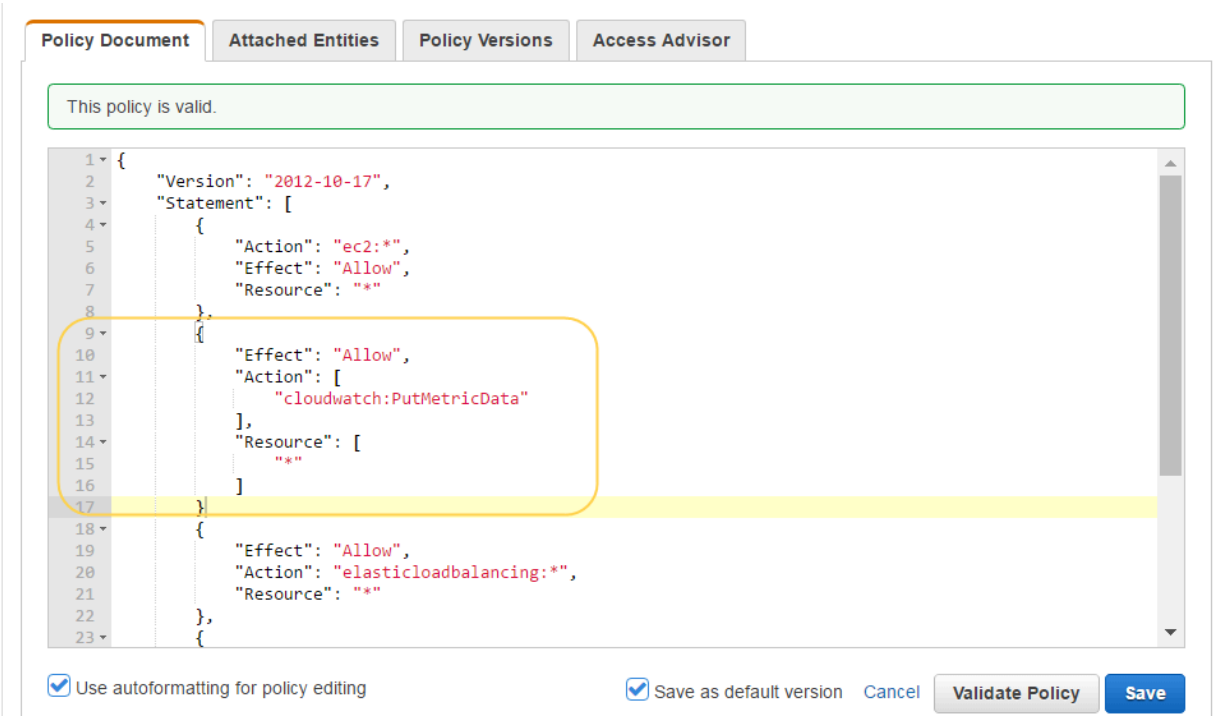

You can copy and the paste the permissions here:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "cloudwatch:PutMetricData"
            ],
            "Resource": [
                "*"
            ]
        }
    ]
}
```

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enable CloudWatch on the VM-Series firewall on AWS.

1. Log in to the web interface on the VM-Series firewall
2. Select **Device** > **VM-Series**.
3. In AWS CloudWatch Setup, click **Edit** ( ) and select **Enable CloudWatch Monitoring**.
   1. Enter the **CloudWatch Namespace** to which the firewall can publish metrics. The namespace cannot begin with **AWS**.

      The aggregated metrics for all VM-Series firewall in an HA pair or auto scaling deployment are published to the namespace you entered above. The namespace with the `_dimensions` suffix that is automatically created enables you to filter and view metrics for an specific VM-Series firewall using the hostname or AWS instance ID metadata attached to the firewall.
   2. Set the **Update Interval** to a value between 1-60 minutes. This is the frequency at which the firewall publishes the metrics to CloudWatch. The default is 5 minutes.

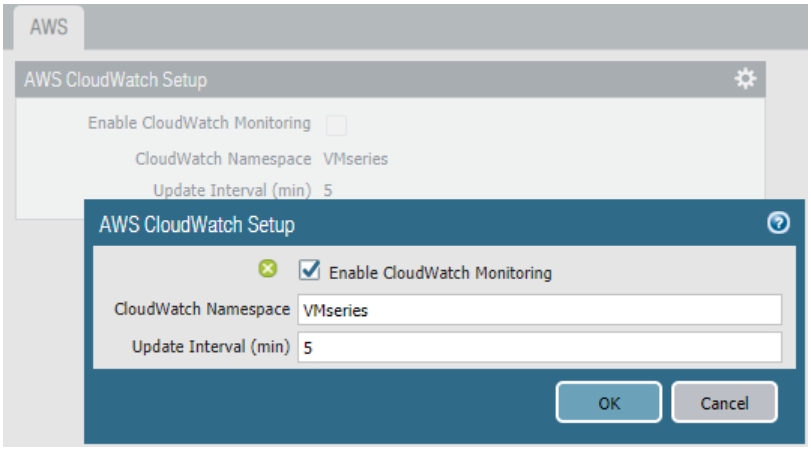

4. **Commit** the changes.

   Until the firewall starts to publish metrics to CloudWatch, you cannot configure alarms for PAN-OS metrics.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Verify that you can see the metrics on CloudWatch.

1. On the AWS console, select **CloudWatch** > **Metrics**, to view CloudWatch metrics by category.
2. From the Custom Metrics drop-down, select the namespace.

   Do not select the default VMseries namespace. This is the seeded data in the VM-Series Firewall. Make sure that you select a different namespace. Also, do not select any namespace that starts with AWS.

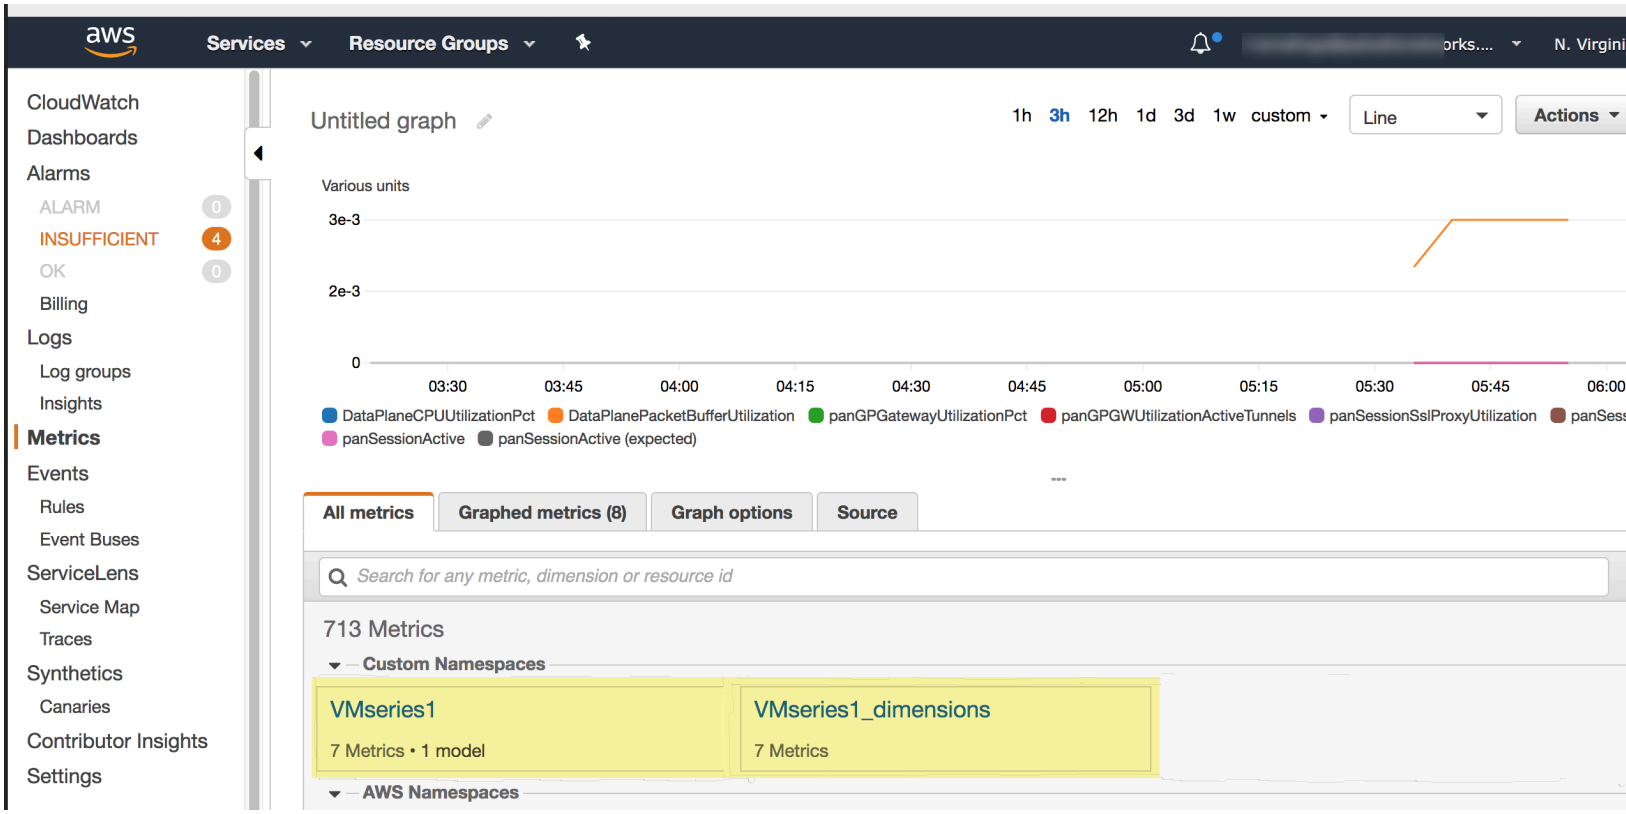

3. Verify that you can see PAN-OS metrics in the viewing list.

   To filter by hostname or AWS Instance ID of a specific firewall, select `_dimensions`.

 ©2024 Palo Alto Networks, Inc.

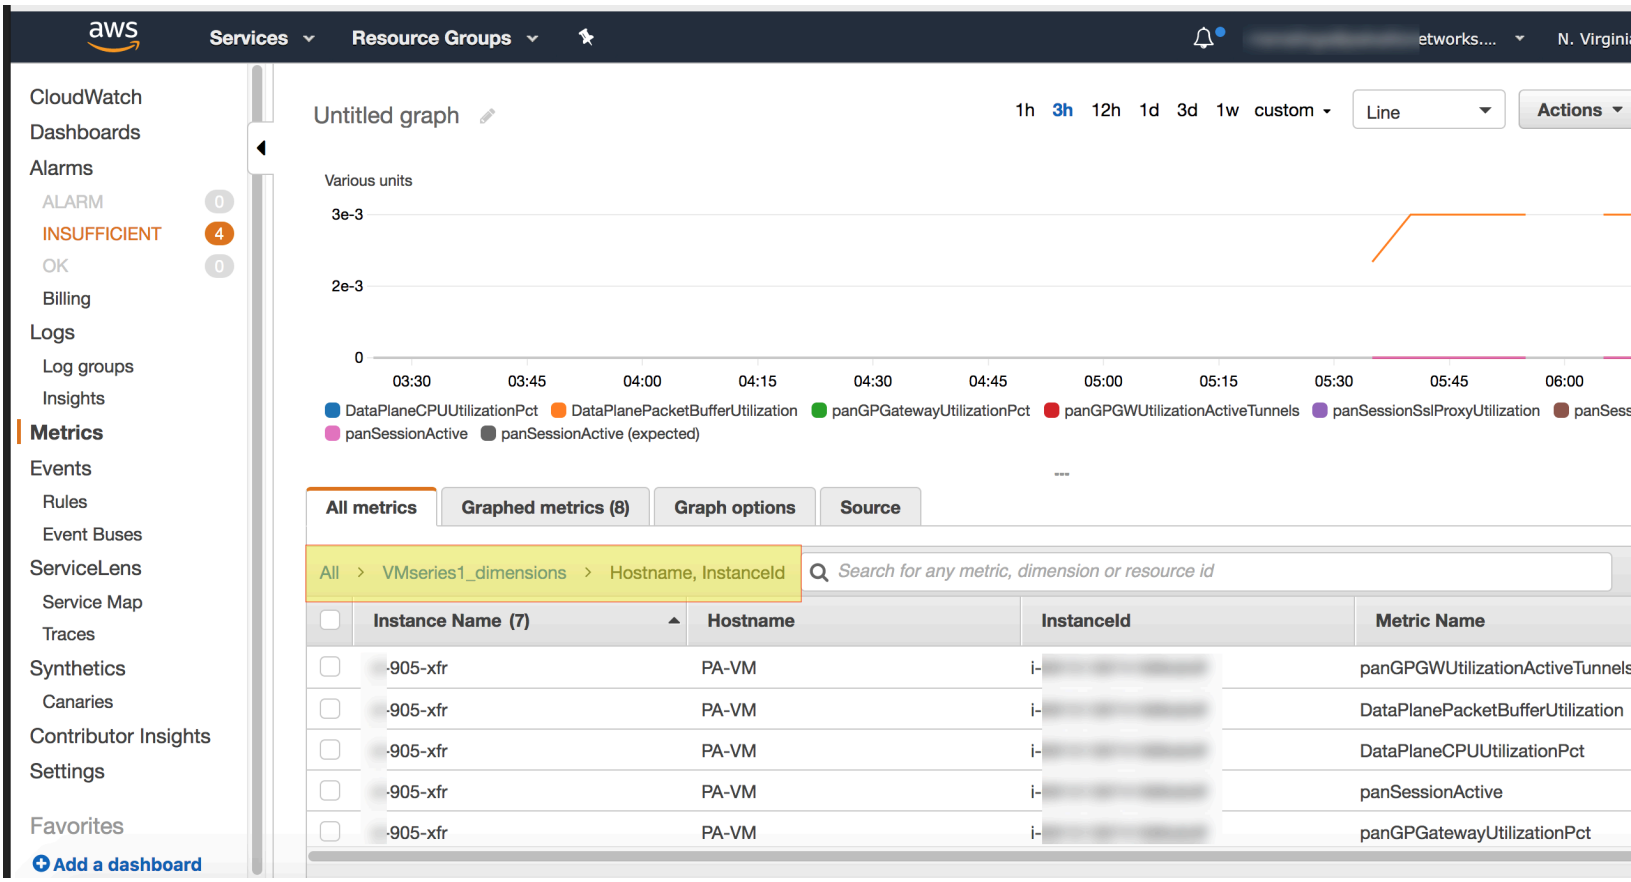

STEP 4 | Configure alarms and action for PAN-OS metrics on CloudWatch.

Refer to the AWS documentation: http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/AlarmThatSendsEmail.html

A VM-Series firewall with bootstrap configuration will take about 7-9 minutes to be available for service. So, here are some examples on how to set alarms that trigger auto scaling for the VM-Series firewall:

- If you have deployed 2 instances of the VM-Series firewalls as Global Protect Gateways that secure remote users, use the GlobalProtect Gateway Active Tunnels metric. You can configure an alarm for when the number of active tunnels is greater than 300 for 15 minutes, you can deploy 2 new instances of the VM-Series firewall, which are bootstrapped and configured to serve as Global Protect Gateways.
- If you are using the firewall to secure your workloads in AWS, use the Session Utilization metric to scale in or scale out the firewall based on resource usage. You can configure an alarm for when the session utilization metric is greater than 60% for 15 minutes, to deploy one instance of the VM-Series instance firewall. And conversely, if Session Utilization is less than 50% for 30 minutes, terminate an instance of the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

Panorama Orchestrated Deployments in AWS

The Panorama Plugin for AWS 3.0.1 or later orchestrates VM-Series firewall deployments in AWS and enables security policies for managed firewalls. AWS Orchestration is designed as a plug and play model for setting up security deployments in AWS. It simplifies the deployment of the existing Gateway Load Balancer (GWLB) solution by bringing all configuration into one screen on Panorama. Panorama lets the plugin manage your deployment and configure resources. This plugin also performs firewall management by generating the needed baseline configuration to get traffic flowing for the deployment. When you configure the policies, the plugin service Inbound, Outbound, and East-West flows for all traffic protocols. Use this plugin to configure, deploy, and manage your security deployments.

The image below highlights the topology of the Security VPC deployment. Here, all security resources are deployed into the plugin managed Security VPC. The GWLB solution is leveraged to redirect traffic from your applications to the firewall stack.

 ©2024 Palo Alto Networks, Inc.

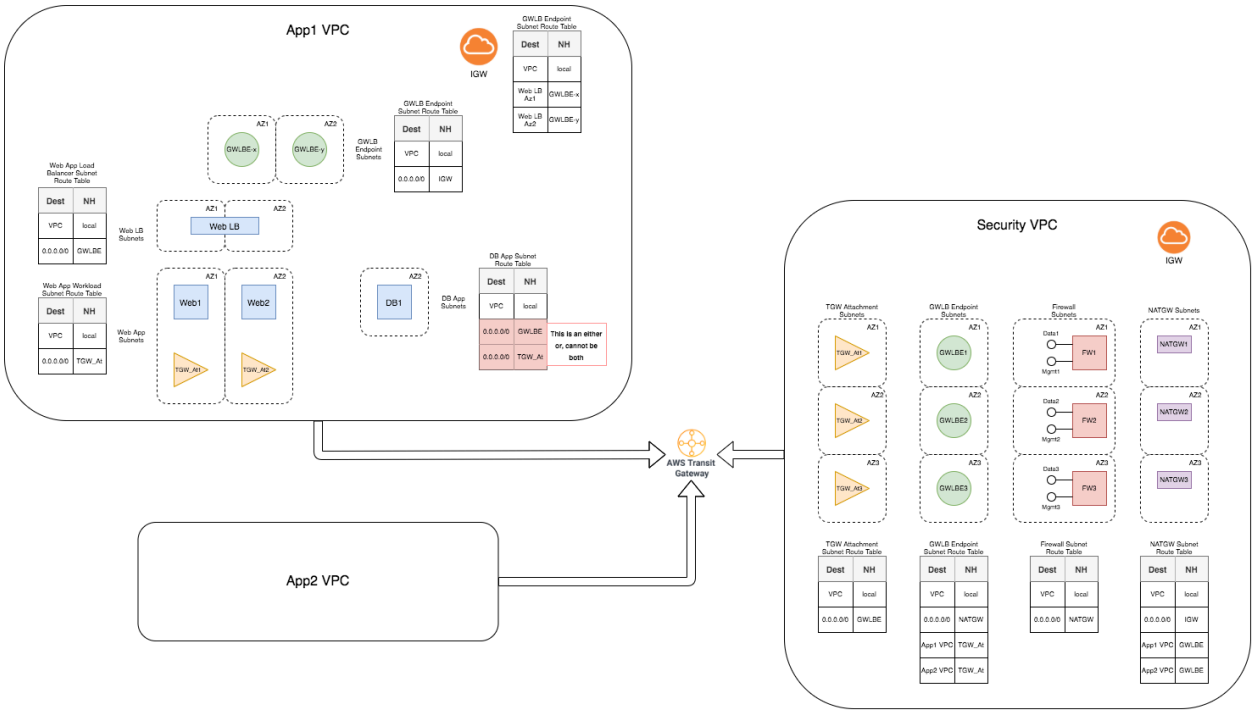

 ©2024 Palo Alto Networks, Inc.

As part of the infrastructure setup on the AWS cloud, the plugin creates Security VPC with GWLB Endpoints, firewalls, and NAT Gateway subnets and route tables. The plugin does not create AWS Transit Gateway (TGW).

VM-Series firewall can inspect traffic routed between the VPCs.

The Inbound traffic flow originating in the Application VPC flows in through IGW is redirected to the GWLB Endpoint based on edge route. The traffic enters through the GWLB Endpoint to the firewalls in the Security VPC for inspection. After the inspection, the traffic is sent back to the GWLB Endpoint and directed to the original application.

For Outbound and East-West traffic, this solution leverages TGW. When you create a TGW, the plugin creates TGW attachments and route tables in the Security VPC. You have to attach your Application VPC to the TGW used in the Security VPC configuration. You must also direct the Outbound and East-West traffic to the TGW by adding routes to the route tables associated with your workload subnets. You have to modify the Application VPC attachment route table to direct the East-West and Outbound traffic to the Security VPC attachment.

The plugin monitors TGW attachments to learn any newly added and deleted VPC attachments. When the plugin detects an existing or new attachment, it makes necessary changes in the Security VPC to ensure that the firewall inspects the traffic entering TGW before sending it back to the TGW. These changes include adding routes to the NAT Gateway route table to direct Outbound traffic back to the GWLB Endpoint, and to GWLB Endpoint route table to return traffic to the TGW after inspection. The plugin updates the TGW attachment route table to ensure that the traffic coming back from the Security VPC to the TGW is sent to the correct Application attachment. Traffic from the Application VPC is directed to TGW through routing. When traffic hits the TGW attachment in Security VPC, the attachment route table sends the traffic to the Security VPC. From there, it is directed to the existing GWLB Endpoint, then to the firewall for inspection. The Outbound traffic flows out to the original destination address through NAT Gateway. The East-West traffic is sent back to the TGW where the route table directs the traffic to the original destination address.

- Prepare for an Orchestrated AWS Deployment
- Orchestrate a VM-Series Firewall Deployment in AWS
- View the Deployment Status
- Traffic Flow and Configurations

Prepare for an Orchestrated AWS Deployment

Complete the following tasks on AWS and Panorama before you orchestrate a VM-Series firewall on AWS.

- **AWS**
  - Panorama-orchestrated deployments on AWS requires two availability zones and supports up to six availability zones.
  - Create an AWS user or instance profile for a specific AWS account with programmatic access and necessary permissions to allow the plugin to create resources on security VPC.

    For Security Account—An AWS Account with an user or instance profile who has the necessary permissions to launch AWS resources such as VPCs, VM instances, AWS load balancer, NAT gateways, and endpoints. The plugin needs the existing user or instance

 ©2024 Palo Alto Networks, Inc.

profile to have the following set of permissions to declare the IAM role valid. The CFT hyperlink under **Security Account** creates a policy with the following permissions.

```
"Statement":[
      {
          "Effect": "Allow",
          "Action": "ec2:*",
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": "elasticloadbalancing:*",
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": "cloudwatch:*",
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": "autoscaling:*",
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": "sts:assumerole",
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": "cloudformation:*",
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": [
              "iam:Get*",
              "iam:List*",
              "iam:PassRole"
          ],
          "Resource": "*"
      },
      {
          "Effect": "Allow",
          "Action": "ram:*",
          "Resource": "*"
      }
    ]
},
```

The following granulized permissions accommodate your requirements and security permissions. These permissions provide a detailed explanation for the API calls made from the plugin. The permissions are granulized to accommodate every action that will be

 ©2024 Palo Alto Networks, Inc.

called from the CFT and plugin back-end code for both Security VPC and cross-account Application VPC.

The following permissions are not implemented in the AWS plugin for Panorama version 3.0.1 because the detailed permissions list exceed the AWS policy size limitation.For inline policies, you can add as many policies as you want for a user, role, or group, but the total aggregate policy size per entity cannot exceed these limits—user policy size cannot exceed 2048 characters, role policy size cannot exceed 10240 characters, and groups policy size cannot exceed 5120 characters. Due to these limitations and back-end validation time, you must use the above mentioned permissions.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "ec2:AuthorizeSecurityGroupIngress",
                "cloudwatch:PutMetricData",
                "ec2:Describe*",
                "cloudwatch:DeleteAlarms",
                "autoscaling:DescribePolicies",
                "ec2:DeleteVpcEndpoints",
                "ec2:AttachInternetGateway",
                "ec2:AcceptTransitGatewayVpcAttachment",
                "autoscaling:ExecutePolicy",
                "ec2:DeleteRouteTable",
                "sts:GetSessionToken",
                "cloudformation:DescribeStackEvents",
                "ec2:RevokeSecurityGroupEgress",
                "ec2:CreateRoute",
                "ec2:CreateInternetGateway",
                "cloudformation:UpdateStack",
                "ec2:DeleteInternetGateway",
                "iam:ListRolePolicies",

 "autoscaling:TerminateInstanceInAutoScalingGroup",
                "iam:ListPolicies",
                "ec2:DisassociateTransitGatewayRouteTable",
                "iam:GetRole",
                "iam:GetPolicy",
                "ec2:CreateTags",
                "elasticloadbalancing:CreateTargetGroup",
                "ec2:RunInstances",
                "ec2:DisassociateRouteTable",
                "ec2:CreateVpcEndpointServiceConfiguration",
                "ec2:CreateTransitGatewayRoute",
                "ec2:CreateTransitGatewayVpcAttachment",
                "elasticloadbalancing:DescribeAccountLimits",
                "elasticloadbalancing:AddTags",
                "cloudformation:DeleteStack",
                "cloudwatch:DescribeAlarms",
```

 ©2024 Palo Alto Networks, Inc.

```
                "ec2:DeleteNatGateway",
                "ram:AssociateResourceShare",
                "autoscaling:DeleteAutoScalingGroup",
                "ec2:CreateSubnet",

"elasticloadbalancing:ModifyLoadBalancerAttributes",
                "iam:GetRolePolicy",
                "ec2:ModifyVpcEndpoint",
                "ec2:DisassociateAddress",
                "autoscaling:DescribeAutoScalingInstances",
                "ec2:ModifyVpcEndpointServicePermissions",
                "ec2:CreateNatGateway",
                "ec2:CreateVpc",
                "ec2:ModifySubnetAttribute",
                "iam:PassRole",
                "autoscaling:DescribeScalingActivities",
                "sts:DecodeAuthorizationMessage",
                "autoscaling:DescribeLoadBalancerTargetGroups",
                "iam:ListAttachedGroupPolicies",
                "ec2:DeleteLaunchTemplateVersions",
                "sts:GetServiceBearerToken",
                "iam:ListAccessKeys",
                "ram:DisassociateResourceShare",
                "ec2:ReleaseAddress",
                "ec2:DeleteLaunchTemplate",
                "elasticloadbalancing:CreateLoadBalancer",
                "ec2:AcceptVpcEndpointConnections",
                "iam:ListGroupPolicies",
                "iam:ListRoles",
                "elasticloadbalancing:DeleteTargetGroup",
                "ram:AssociateResourceSharePermission",
                "ec2:CreateLaunchTemplate",
                "elasticloadbalancing:DescribeTargetGroups",
                "elasticloadbalancing:DeleteListener",
                "ram:UpdateResourceShare",
                "iam:GetPolicyVersion",
                "ec2:DeleteSubnet",
                "ec2:ModifyVpcEndpointServiceConfiguration",
                "ec2:CreateTransitGatewayRouteTable",
                "ec2:ModifyTransitGateway",
                "cloudformation:DescribeStackResource",
                "ec2:AssociateRouteTable",
                "elasticloadbalancing:DeleteLoadBalancer",
                "elasticloadbalancing:DescribeLoadBalancers",
                "logs:CreateLogStream",
                "ec2:GetLaunchTemplateData",
                "ec2:DeleteTransitGatewayVpcAttachment",
                "autoscaling:DescribeAutoScalingGroups",
                "iam:ListAttachedRolePolicies",
                "logs:GetLogEvents",
                "autoscaling:UpdateAutoScalingGroup",
                "ec2:AssociateTransitGatewayRouteTable",

"elasticloadbalancing:ModifyTargetGroupAttributes",
                "autoscaling:SetDesiredCapacity",
                "cloudformation:DescribeStackResources",
```

 ©2024 Palo Alto Networks, Inc.

```
                "ec2:CreateRouteTable",
                "ec2:DetachInternetGateway",
                "cloudformation:DescribeStacks",
                "ec2:DeleteTransitGatewayRouteTable",
                "sts:AssumeRole",
                "ec2:DeleteTransitGatewayRoute",
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "ec2:DeleteVpc",
                "iam:GetGroupPolicy",
                "ec2:AssociateAddress",
                "autoscaling:CreateAutoScalingGroup",
                "ram:AcceptResourceShareInvitation",
                "ec2:DeleteTags",
                "logs:DescribeLogStreams",
                "ec2:DeleteVpcEndpointServiceConfigurations",
                "autoscaling:DeletePolicy",
                "elasticloadbalancing:RemoveTags",
                "elasticloadbalancing:CreateListener",
                "elasticloadbalancing:DescribeListeners",
                "autoscaling:PutScalingPolicy",
                "ec2:CreateSecurityGroup",
                "iam:ListAttachedUserPolicies",
                "ec2:ModifyVpcAttribute",
                "ec2:ModifyInstanceAttribute",
                "ec2:GetTransitGatewayRouteTableAssociations",
                "ram:DeleteResourceShare",
                "ec2:AuthorizeSecurityGroupEgress",
                "ec2:ModifyTransitGatewayVpcAttachment",
                "iam:GetInstanceProfile",
                "ram:DisassociateResourceSharePermission",
                "elasticloadbalancing:DescribeTags",
                "ec2:DeleteRoute",
                "iam:ListUserPolicies",
                "logs:PutLogEvents",
                "ec2:AllocateAddress",
                "ec2:CreateLaunchTemplateVersion",
                "cloudwatch:PutMetricAlarm",
                "cloudformation:CreateStack",
                "ec2:CreateVpcEndpoint",
                "ec2:DeleteSecurityGroup",

"ec2:StartVpcEndpointServicePrivateDnsVerification",
                "ec2:ModifyLaunchTemplate",
                "iam:ListUsers",
                "ram:CreateResourceShare"
            ],
            "Resource": "*"
        }
    ]
```

 ©2024 Palo Alto Networks, Inc.

```
}
```

For Application Account—An AWS account other than the Security account that hosts either TGW or the applications that needs to be protected. Within this account you must create a RoleARN with the following permissions.

```
{
          "Version": "2012-10-17",
          "Statement": [
                {
                    "Action": [
                        "ec2:Describe*",

 "ec2:DisassociateTransitGatewayRouteTable",
                        "ec2:CreateTransitGatewayRoute",
                        "ec2:CreateTransitGatewayRouteTable",
                        "ec2:AssociateTransitGatewayRouteTable",
                        "ec2:DeleteTransitGatewayRouteTable",
                        "ec2:DeleteTransitGatewayRoute",

 "ec2:GetTransitGatewayRouteTableAssociations"
                    ],
                    "Resource": "*",
                    "Effect": "Allow"
                },
                {
                    "Action": [
                        "iam:Get*",
                        "iam:List*"
                    ],
                    "Resource": "*",
                    "Effect": "Allow"
                }
          ]
}
```

- **Dedicated CIDR block**—A CIDR block reserved for the Security VPC. The plugin manages this CIDR block, using it to launch firewalls, load balancers, and other deployment resources for the Security VPC.
- **AWS transit gateway**—Create a TGW and ensure that the selected AWS user has permission to configure the TGW resources.

- **Panorama**
  - **Panorama Plugin for AWS**—Version 3.0.1 or later.
  - **VM-Series Plugin**—Version 2.0.6 or later.
  - **PanOS**—Version 10.0.5 or later.
  - Create a valid license API key configured on Panorama for delicensing the firewalls.
  - Create an IAM role on the plugin under **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Roles**. This configuration needs the Access Key and Secret Key associated with the user you created in your AWS account.

 ©2024 Palo Alto Networks, Inc.

Configure IAM Roles for AWS Plugin in Panorama

With the AWS plugin 3.0.1 or later, you can use IAM roles to enable Panorama to authenticate and retrieve metadata on the resources deployed within your AWS account(s). When your Panorama is not deployed on AWS, you have two options. You can either provide the long-term IAM credentials for the AWS accounts, or set up an Assume Role on AWS to allow access to the defined AWS resources within the same AWS account or cross-accounts. An Assume Role is recommended as the more secure option.

STEP 1 | To validate the AWS user credentials created for Security VPC, go to **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Roles**.

STEP 2 | Click **Add** and enter the following details under **Security Account Detail**.

- Enter a name for the IAM role and an optional description.
- Enter the AWS access key and secret key to validate permissions. Re-enter the secret key to confirm the secret access key.
- Select an account type—**Instance Profile** or **AWS Account Credentials**. If your Panorama is deployed on AWS, you can choose to either attach an instance profile with the correct permissions to your Panorama or add the credentials associated with the IAM role on Panorama. If your Panorama is not deployed on AWS, you must enter the credentials for the IAM role locally on Panorama.

STEP 3 | Under **Application Account Details**, search and select the needed RoleARNs to provide valid permissions to the Security account to access the resources in the Application VPC.

The status of validity of monitoring and deployment is color coded for ease of identification.

- Valid (Green)—Indicates that the secret key and access key are valid. Also, all RoleARNs entered for application account access have valid permissions to do necessary action.
- Partially valid (Orange)—Indicates that the secret key and access key are valid but one or more RoleARNs entered for application account access do not have valid permissions to do necessary action. Click the status hyperlink to open the IAM and see which specific RoleARNs do not comply.
- Invalid (Red)—Indicates that the secret key and access key entered are either invalid or do not have permissions to do the necessary action.
- Commit Required (Gray)—Indicates that a commit is required for the role.
- Validating (Gray)—Indicates that the plugin is trying to connect to AWS to check the necessary requirements. If this status continues for more than a few seconds, verify if the connection to AWS is established.

Only the IAM roles with green or orange status are allowed for further deployment configuration.

Orchestrate a VM-Series Firewall Deployment in AWS

Complete the following procedure to orchestrate a VM-Series firewall deployment in AWS.

Panorama deployments on AWS with an instance profile is not supported when deployed behind a proxy.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Log on to the Panorama web interface.

STEP 2 | Install Panorama plugin for AWS 3.0.1 or later.

To upgrade the Panorama plugin for AWS to version 3.0.1, you must first upgrade the plugin to version 2.0.2. After you install the AWS plugin version 3.0.1 you cannot downgrade to version 2.0.x or below.

If you have a Panorama HA configuration, repeat the installation/upgrade process on each Panorama peer.

If you currently have a Panorama plugin for any cloud platform installed, installing (or uninstalling) an additional plugin requires a Panorama reboot so that you can commit changes.

STEP 3 | Configure IAM Roles for AWS Plugin in Panorama

STEP 4 | Select **Panorama** > **Plugins** > **AWS** > **Deployments** to **Add** a new deployment.

STEP 5 | Enter the generic details of the deployment in the **General** tab.

- Enter a **Name** and an optional **Description** to identify the deployment in Panorama and AWS cloud.
- Select an **IAM Role** from the drop-down. The list displays IAM roles that has valid or partially valid deployment permissions.

Once an IAM role has been created and added to you deployment, you can edit information in the IAM (such as secret key and access key). However, you cannot edit the name of the IAM role. To change the name, you must delete the IAM and create another.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Enter the Security VPC related information in the **Security VPC** tab.

- Select the AWS **Region** in which you intend to launch the deployment. The list displays regions based on the selected IAM role.
- Enter a **VPC CIDR** value to create resources in the Security VPC. This CIDR will be managed by the AWS plugin.
- Select two or more **Availability Zones** from the pre-populated list and follow the same mapping in AWS. This list is populated based on the region you selected.

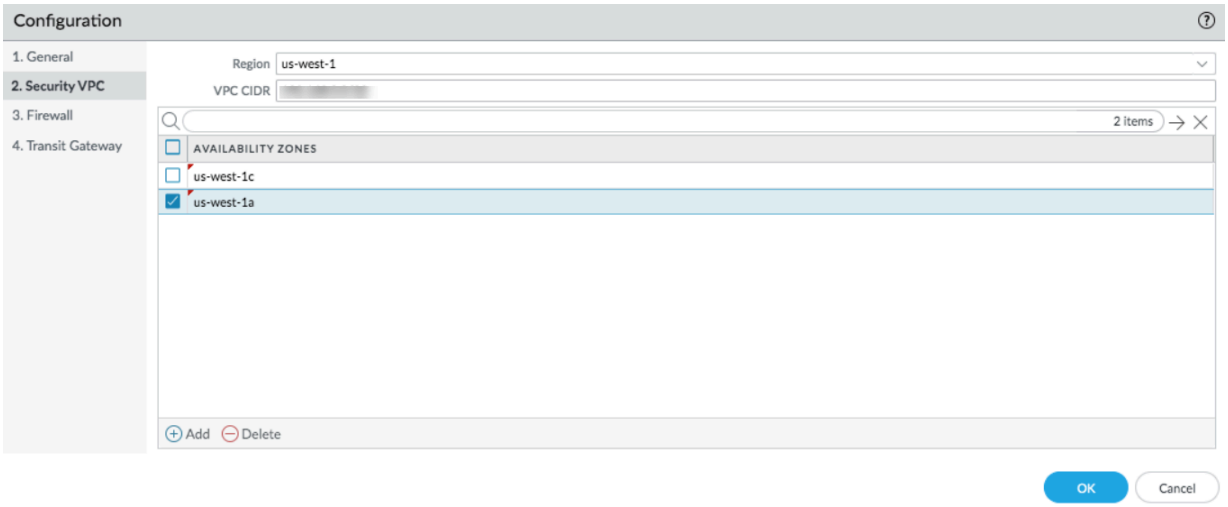

STEP 7 | Select **Firewall** > **Image** and enter the following details.

- **License Type**—The standard license types **Bring Your Own License** (BYOL), **Pay As you Go-Marketplace-Bundle1**, and **Pay As you Go-Marketplace-Bundle2** are provided as

 ©2024 Palo Alto Networks, Inc.

options in the drop-down (based on the selected regions). If you select **Bring Your Own License**, be prepared to enter a license authcode.

Pay as you go Bundle 1 and Bundle 2 cannot be used with an AMI custom image.

- (Optional—Appears only if you choose **Bring Your Own License** license type) **License Authcode**—Enter the authcode for your BYOL. This authcode determines which instance types appear on the **Instance Type** drop-down.
- **Instance Type**—Choose supported instance types from the drop-down. This list is derived from the license authcode.
- **Image Type**—Select Marketplace Image or Custom Image.

  If you select **Marketplace Image**, select from the drop-down, **PanOS Version** 10.0.5 or later supported in the regions you selected when you configured the Security VPC.

  If you select **Custom Image**, enter the Amazon Machine Image (AMI) ID and select a PanOS Version 10.0.5 or later.
- **Device Certificate**—The device certificate is generated on the Customer Support portal, and enables you to retrieve your site license entitlements for AutoFocus or Cortex Data Link. Select **Disable** if you are not using these licenses. To configure the device certificate PIN, select **Enable** and enter the following information:

  PIN ID—Enter the PIN ID.

  Confirm PIN ID—Re-enter the PIN ID.

  PIN Value—Enter the PIN.

  Confirm VM PIN Value—Re-enter the PIN.

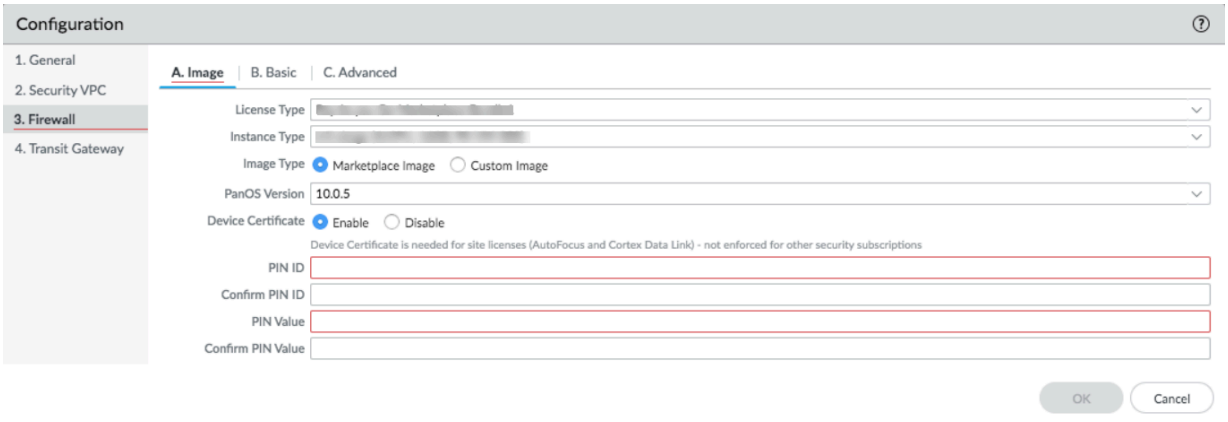

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Select **Firewall** > **Basic** and enter the following details.

- **AWS Key Name**—The name of a SSH key you will use to log into the firewalls after they are deployed. This key is bootstrapped into the firewall and can be used for debugging when the firewall is up and running.
- **Existing Device Group**—If you select **No**, the plugin creates format of the device group name. If you select **Yes**, select an existing **Device Group** from the drop-down list.
- **Primary Panorama IP**—The IP address of the Panorama you are using. The drop-down displays public and private IP addresses on the management interface. Select an IP address from the drop-down.

  *If you deployed Panorama behind a proxy, you must manually enter the public IP of the primary Panorama under **Panorama** > **Setup** > **Interfaces**.*
- **Secondary Panorama IP**—If you have a Panorama HA, the drop-down displays the IP addresses on the management interface of the secondary device. Select an IP address from the drop-down.

  If the secondary Panorama has a public IP address, it may not appear in the drop-down. In this case, you must manually add the IP address of the secondary Panorama.

  If Panorama private IP is used, routes may need to be added to the AWS route tables for plugin deployed firewall subnets and existing Panorama subnet, in order to facilitate connectivity between the newly deployed firewalls and Panorama.
- **Min Firewalls**—The minimum number of firewalls in an Auto Scaling Group (ASG). A value between 1 and 25.
- **Max Firewalls**—The maximum number of firewalls in an ASG. A value between 2 and 25.
- **FirewallInstanceARN**—From the drop-down, choose the assume RoleARN created on AWS cloud that is associated with the firewall instance to publish autoscaling metrics. The drop-down displays only the RoleARNs you entered on the **Setup** > **IAM Roles** page.

 ©2024 Palo Alto Networks, Inc.

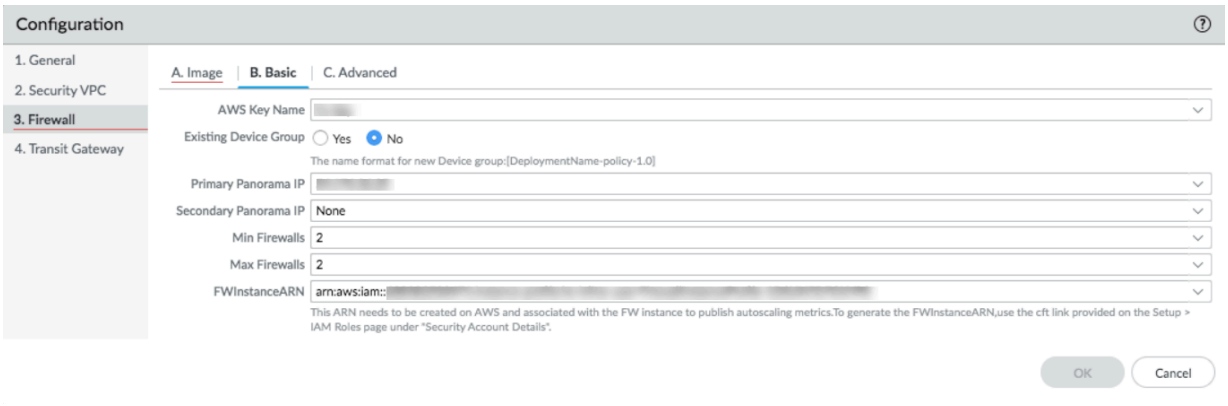

STEP 9 | (Optional) Select **Firewall** > **Advanced** and enter the following details.

- **Autoscaling Metric**—Choose a metric from the drop-down: Data Plane CPU Util Percent (default), Active Sessions, Data Plane Packet Buffer Util Percent, or Session Util Percent.
- **Scale In Threshold**—Choose a value for the scale in threshold. The value depends on your chosen metric.
- **Scale Out Threshold**—Choose a value for the scale out threshold. The value depends on your chosen metric.
- **Scale Out Threshold**—Choose a value for the scale out threshold. The value depends on your chosen metric.
- **Jumbo Frame**—Disabled by default. You can only enable this option when preparing the initial deployment. Select **Enable** to enable jumbo frame support on the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Select whether to connect to a **Transit Gateway** to handle traffic routing across Security VPC and Application VPC.

- Choose if you want to **Connect to TGW**. If you select **Yes**, be prepared to enter a TGW ID to which you want to attach the Security VPC.

  This configuration is required for Outbound or East-West traffic flows only.

- (Optional) Select a **TGW ID** to which you want to attach the Security VPC.

  *You have to share the TGW if you want to use them across accounts. You can share it using **Resource Access Manager** (RAM) on AWS. Create RAM based on the account where the TGW is located.*

- Select **Application Account Names**. If the TGW and Security VPC are in the same account, select the Application Account with which you want to share the TGW. The plugin creates the RAM on the Security Account to share the TGW across the selected

 ©2024 Palo Alto Networks, Inc.

Application accounts. You must accept the invitation for RAM on the account you select here.

If the TGW and Security VPC are in the same account, select an Application account with which you want to share the TGW. If the TGW is in an Application account, make sure that the TGW is shared on RAM.

If the TGW is in an Application account (other than the Security account):

1. Make sure the TGW is shared with the Security account.
2. Use the CFT hyperlink under **Setup** > **IAM Roles** > **Application Account Details**. From the CFT, you can create the RAM for the mentioned TGW.
3. On the Security account, make sure to go to RAM in the AWS console and accept the request to share the TGW.

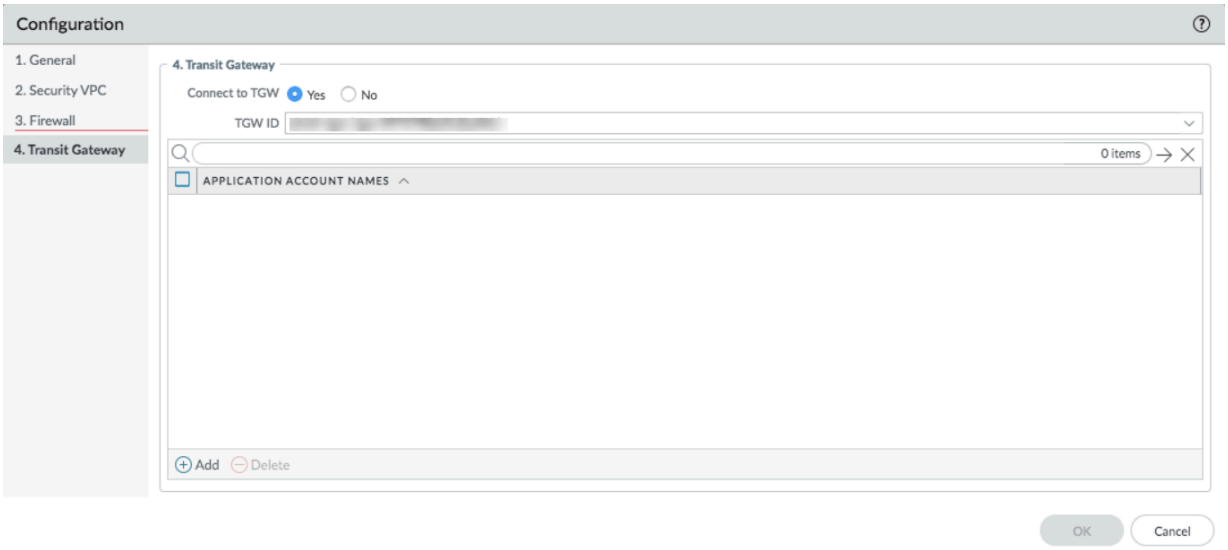

STEP 11 | **Commit** to add the deployment and push to firewalls.

View the Deployment Status

If there is an entry in the **Deployment Status** column, click the hyperlink to view the deployment details.

 ©2024 Palo Alto Networks, Inc.

The possible status messages are:

- **Commit changes**—You have added a deployment for the first time but have not yet committed the changes.

  Every configuration change for the deployment must be committed so that the plugin can pick up your changes.

- **Deploying**—The plugin is deploying or updating the deployment. For more information, click the hyperlink to view the detailed status.
- **Failure**—Deployment has failed. Click the hyperlink and view the **Detailed Status** for the Security stack.
- **Not Deployed**—The plugin is ready to deploy the configuration, but the deployment has not begun.
- **Success**—The plugin has successfully deployed the Security stack and the firewalls have connected to Panorama. The firewalls can pass traffic.
- **Warning**—Deployment has successfully finished but something external to the deployment has failed. For example, you might see this message:

  `FWs have not connected after 20 minutes of the deployment completing.`

  Click the hyperlink and view the Security stack.

Once the deployment is deployed, the plugin allows you to modify a certain subset of parameters. Once the changes have been made, you must do a commit before clicking the **Redeploy** button. When an update happens, the plugin makes sure the Panorama config is created and accurate. It redeploys the CFT to apply any changes, and attach or detach from the configured TGW (if this configuration was modified).

- **Deploy**—After you commit your initial configuration, select **Deploy** to launch the deployment.
- **Redeploy**—Modify a deployment, commit your changes, and select **Redeploy**.

  *You must commit changes to the deployment before you click **Redeploy**.*

- **Undeploy**—Delete a deployment, but keep the configuration so it can be redeployed at a later time.

*To remove an existing deployment and its configuration, check a deployment and select **Delete** at the bottom of the **Deployments** page.*

Detailed Status

To access the **Detailed Status**, click the hyperlink in the **Deployment Status** column. From the detailed status you can learn where to apply configuration, view the error message from a stack failure, or view the deployment status when it is deploying.

- **Name**—The deployment name.
- **Status**—See Deployment Status for description of each status.
- **Detail**—Details on the deployment you selected in **Deployment Status**. For example, if the deployment was successful, displays the date and time of the deployment, or if there was a stack failure, displays an error message.

 ©2024 Palo Alto Networks, Inc.

- **Policy Device Group**—The plugin can create a policy device group for your deployment or you can choose an existing device group to act as the policy device group for a specific deployment.
- **Config Device Group**—The plugin creates a configuration device group as a child of the policy device group. The plugin puts configuration information for the deployment in the config device group, ensuring that your policy device group remains untouched if you remove the deployment.

  Do not put policy information in the config device group.

- **Template Stack**—Displays the template stack associated with the VM-Series firewall. Any custom configuration is applied to this template stack.
- **External IP**—Displays the public IP addresses of the NAT Gateways in the Security VPC, one for each availability zone. The outbound public IP addresses are used for all outbound traffic from the deployment, and for outbound traffic from the VM-Series firewall management interface.

  To allow firewalls to connect to Panorama, the outbound public IP addresses must be whitelisted in your Panorama security group.

- **CloudFormation Link**—This link opens the AWS console to display the current stack in the Cloud Formation services section. You can see where the stack is deployed and debug issues with the deployment.
- **CloudWatch Link**—This link opens the AWS console to display **PaloAltoNetworkFirewalls** logs and log groups related to the firewall.
- **AutoScalingGroup Link**—This link opens the AWS console to display the details of the ASG associated with the deployment, and list of instances under the ASG. You can view logs associated with these instances on **CloudWatch Link**.
- **Endpoint Service Name**—The GWLB Endpoint name created as part of the deployment. For example, com.amazonaws.vpce.us-east-1.vpce-svc-0d00ebcb0000dc000.
- **Cloudformation Stack Name**— For example mynw-aws2-virgexstdg0-c0b0f.

Traffic Flow and Configurations

The plugin deploys and manages the Security VPC. The plugin updates the Security VPC route tables based on the attachments discovered on the AWS Transit Gateway.

Inbound Traffic Flow

Table 2: Inbound traffic flow combinations

	Application	Traffic Type
1	In Security Account	Inbound
2	In Application Account	Cross-Outbound

Use Case: Inbound Traffic - Application is in the Security Account

 ©2024 Palo Alto Networks, Inc.

The plugin creates a VPC Service Endpoint on the Security Account. The GWLB Endpoints must be associated with the VPC Endpoint Service.

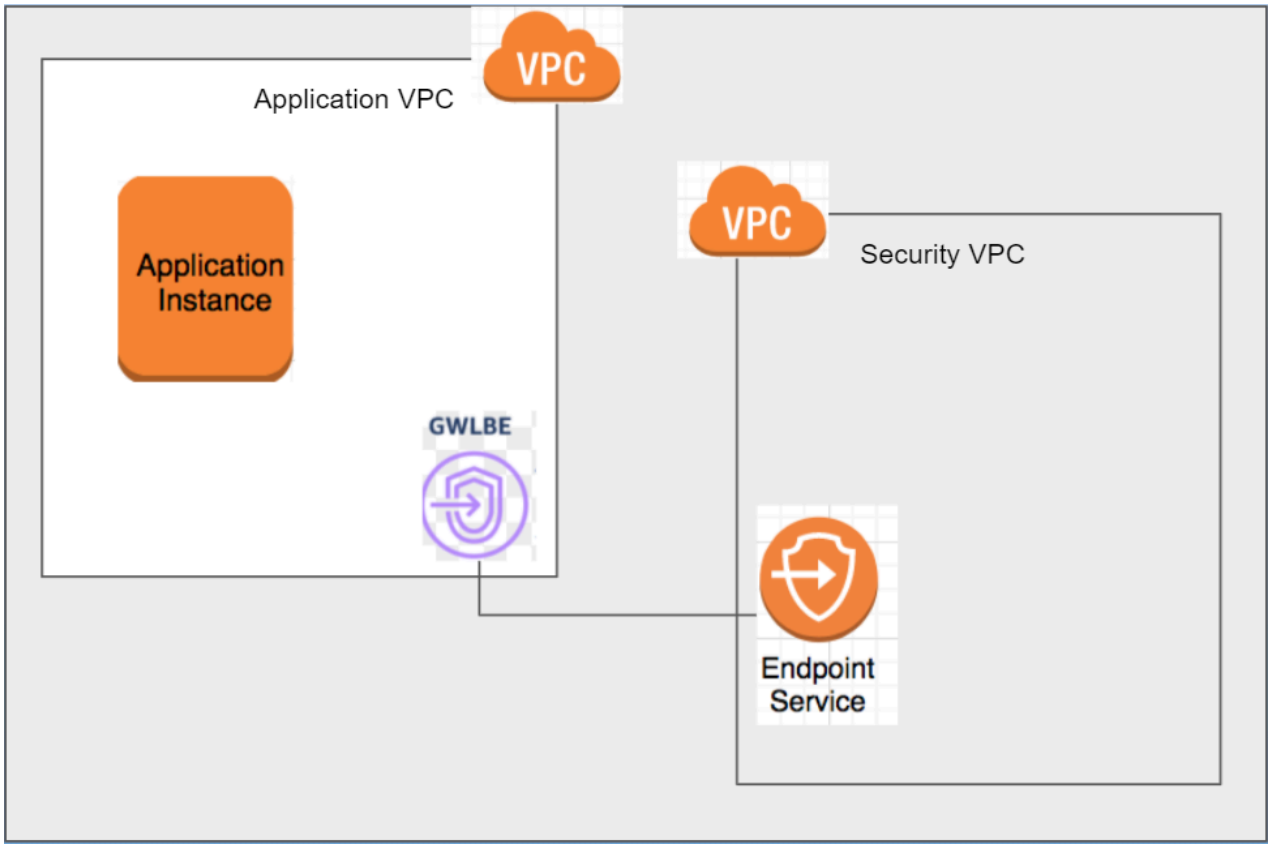

Use Case: Inbound Traffic - Application is in other Application Account

When the application is in a different account, on the AWS console in the navigation pane, choose **Endpoint Services** and select your Endpoint Service. Select **Actions** > **Add Principal** to allow principals. For example, **`arn:aws:iam::AccountNumber:root`**. The GWLB Endpoints must be associated with the VPC Endpoint Service.

 ©2024 Palo Alto Networks, Inc.

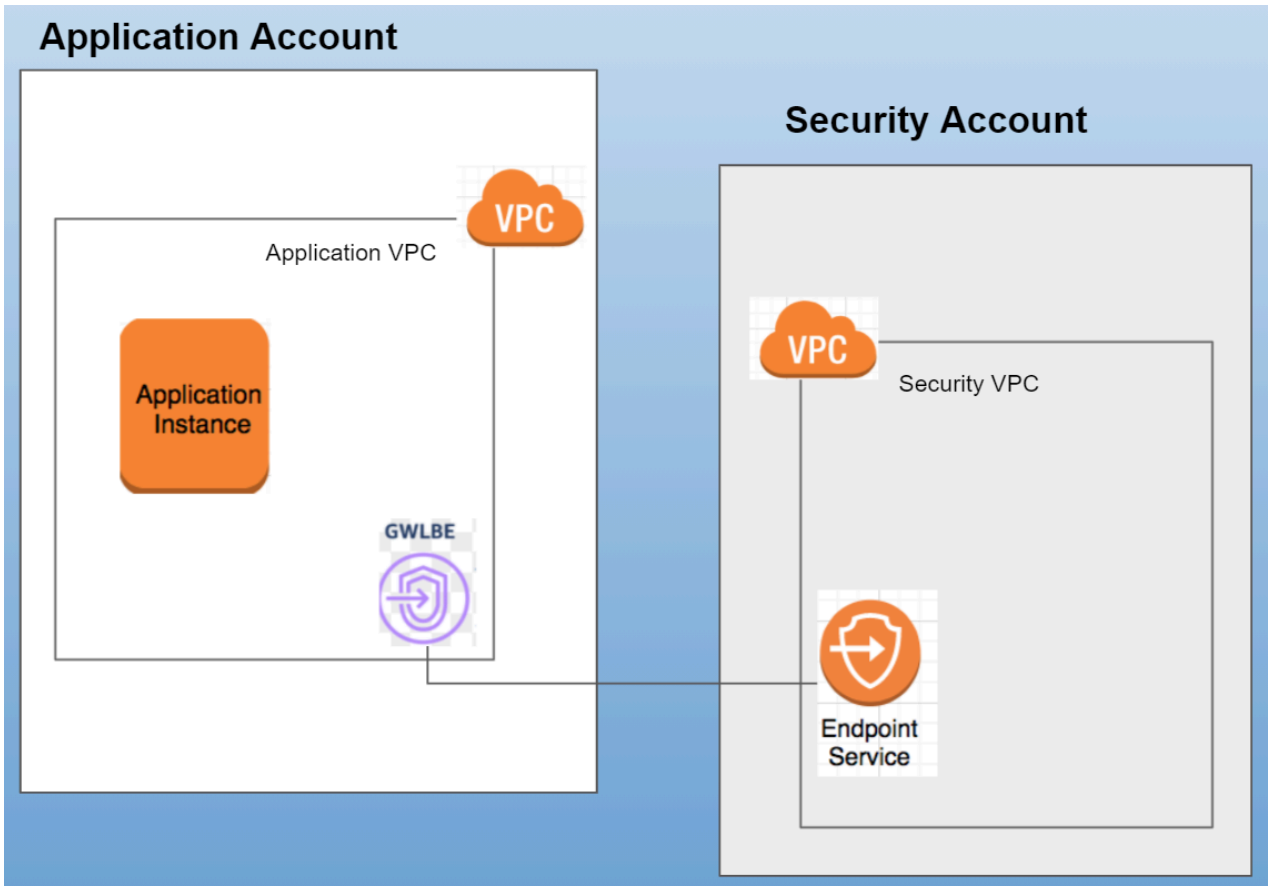

Outbound and East-West Traffic Flow

Table 3: Outbound traffic flow combinations

	Transit Gateway	Application	Traffic Type
1	In Security Account	In Security Account	Outbound
2	In Security Account	In Application Account	Outbound

 ©2024 Palo Alto Networks, Inc.

	Transit Gateway	Application	Traffic Type
3	In Application Account	In Application Account	Cross-Outbound
4	In Application Account	In Security Account	Cross-Outbound

Use Case: Outbound Traffic - Transit Gateway and Application is in the Security Account

The plugin scan for the attachments on the configured TGW. When the plugin detects an existing or new attachment, it makes necessary route table modifications on the Security VPC components.

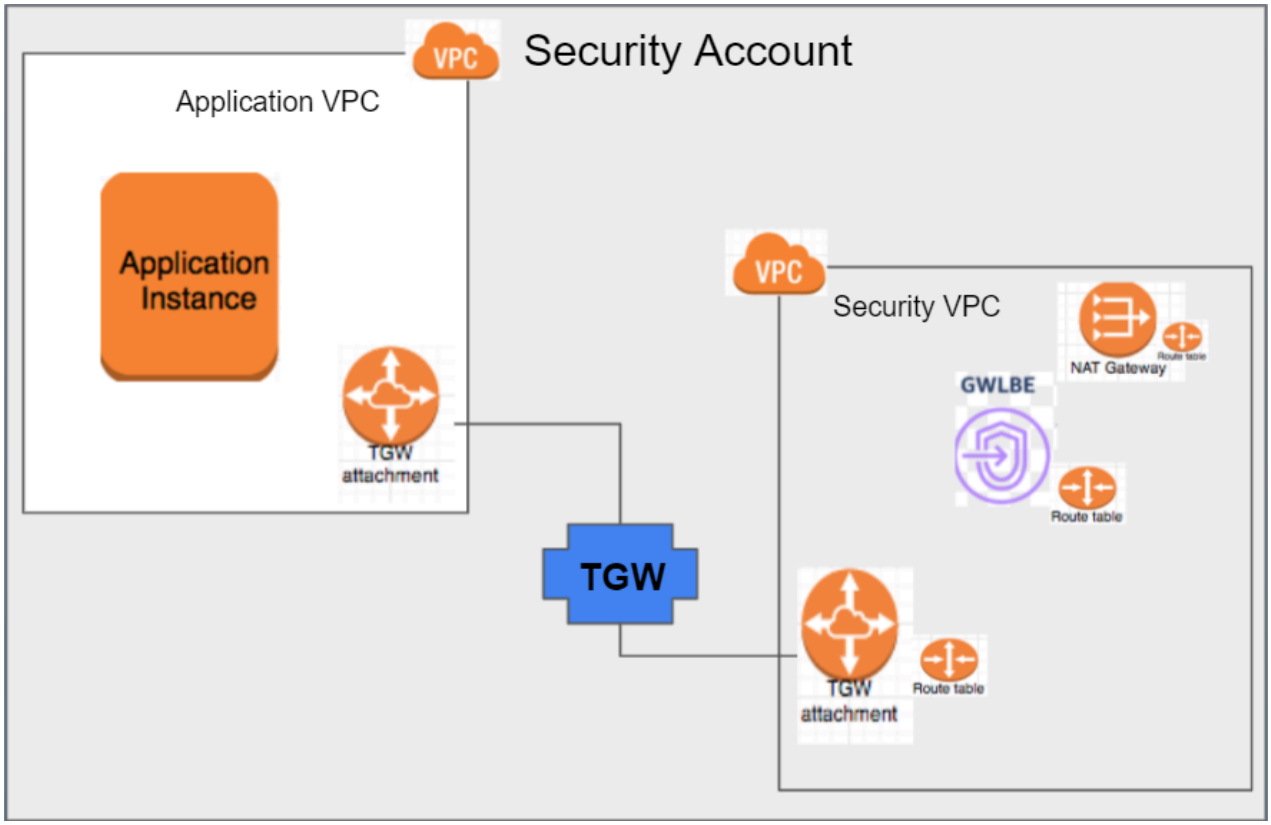

Use Case: Outbound Traffic - Transit Gateway is in Security Account and Application is in the Application Account

When TGW is in the Security Account, to protect the applications that are not in the Security Account, the TGW is shared across these applications using Resource Access Manager (RAM) in the AWS console. You can choose the accounts with which you want to share the TGW from the plugin user interface. Once the deployment is in **Deploying** state, monitor the RAM on the Application Account for an invitation to share resources.

 ©2024 Palo Alto Networks, Inc.

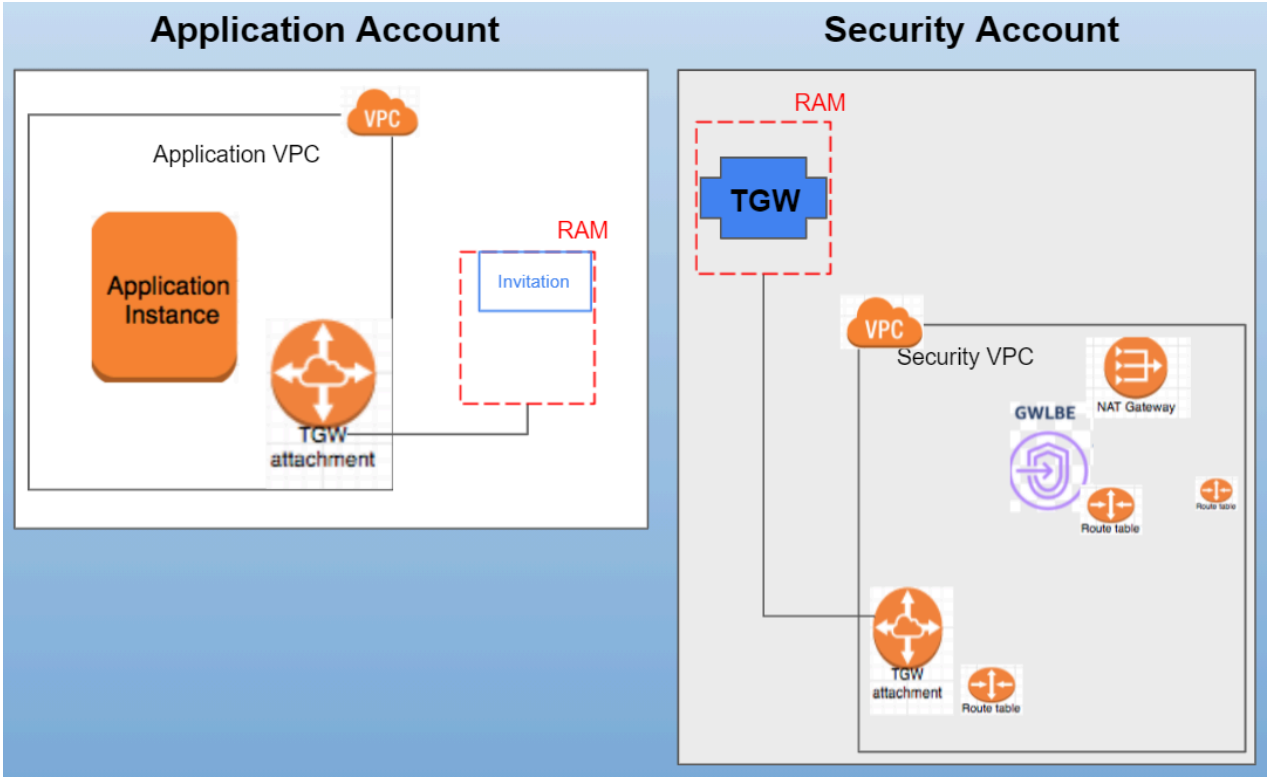

Use Case: Outbound Traffic - Transit Gateway and Application are in the Application Account

When TGW is the Application Account, it must be shared with the Security Account using the RAM. To create a TGW attachment and route table, a RoleARN from this account must be added to the IAM role used for the deployment. Use the CFT hyperlink under **Setup** > **Application Account** to configure the Application Account prerequisites.

 ©2024 Palo Alto Networks, Inc.

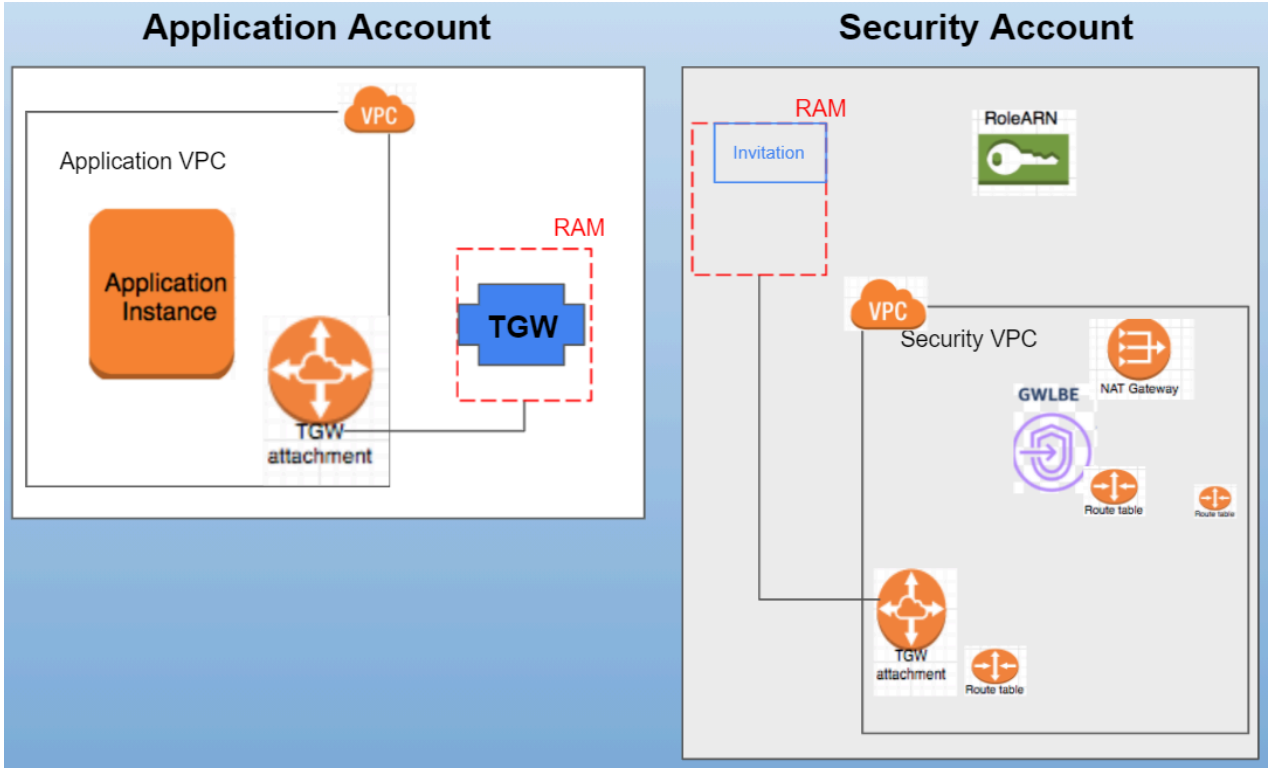

Table 4: East-West traffic flow combinations

	Transit Gateway	Application 1	Application 2	Traffic Type
1	In Security Account	In Security Account	In Security Account	East-West
2 (multi account application)	In Security Account	In Security Account	In Application Account	East-West
3	In Application Account	In Application Account	In Application Account	Cross East-West
4 (multi account application)	In Application Account	In Application Account	In Security Account	Cross East-West

Use Case: East-West Traffic - Transit Gateway and Application1 are in the Security Account and Application2 is in the Security Account

 ©2024 Palo Alto Networks, Inc.

When TGW is in the Security Account, to protect the applications that are not in the Security Account, the TGW is shared across these applications using Resource Access Manager (RAM) in the AWS console. You can choose the accounts with which you want to share the TGW from the plugin user interface. Once the deployment is in **Deploying** state, monitor the RAM on the Application Account for an invitation to share resources.

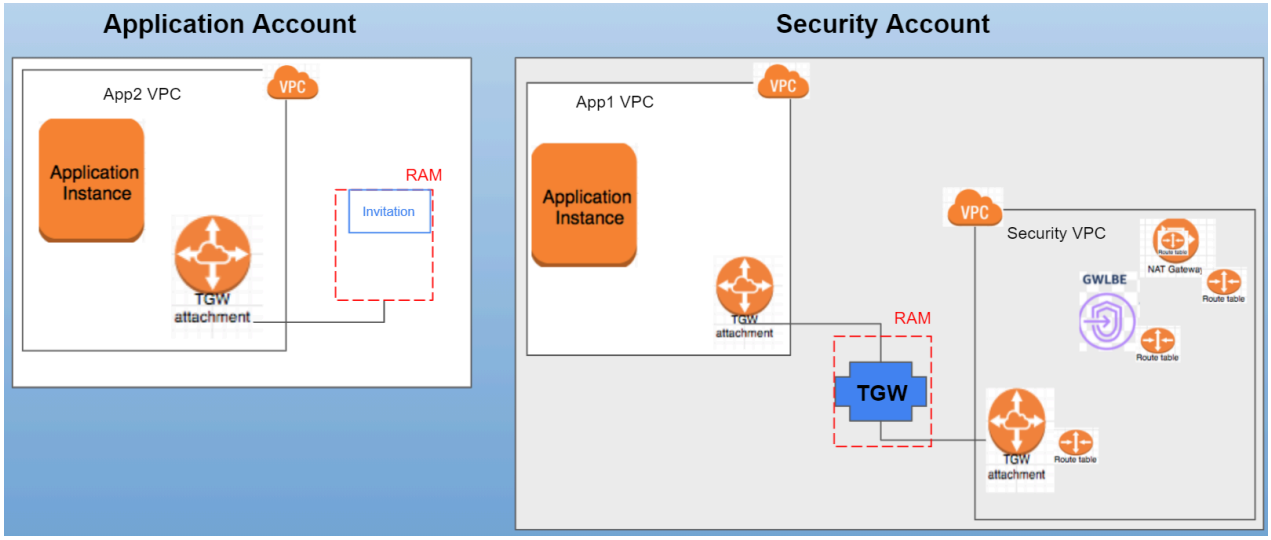

Use Case: East-West Traffic - Transit Gateway and Application1 are in the Application Account and Application2 is in the Security Account

When TGW is the Application Account, it must be shared with the Security Account using the RAM. To create a TGW attachment and route table, a RoleARN from this account must be added to the IAM role used for the deployment. Use the CFT hyperlink under **Setup** > **Application Account** to configure the Application Account prerequisites.

 ©2024 Palo Alto Networks, Inc.

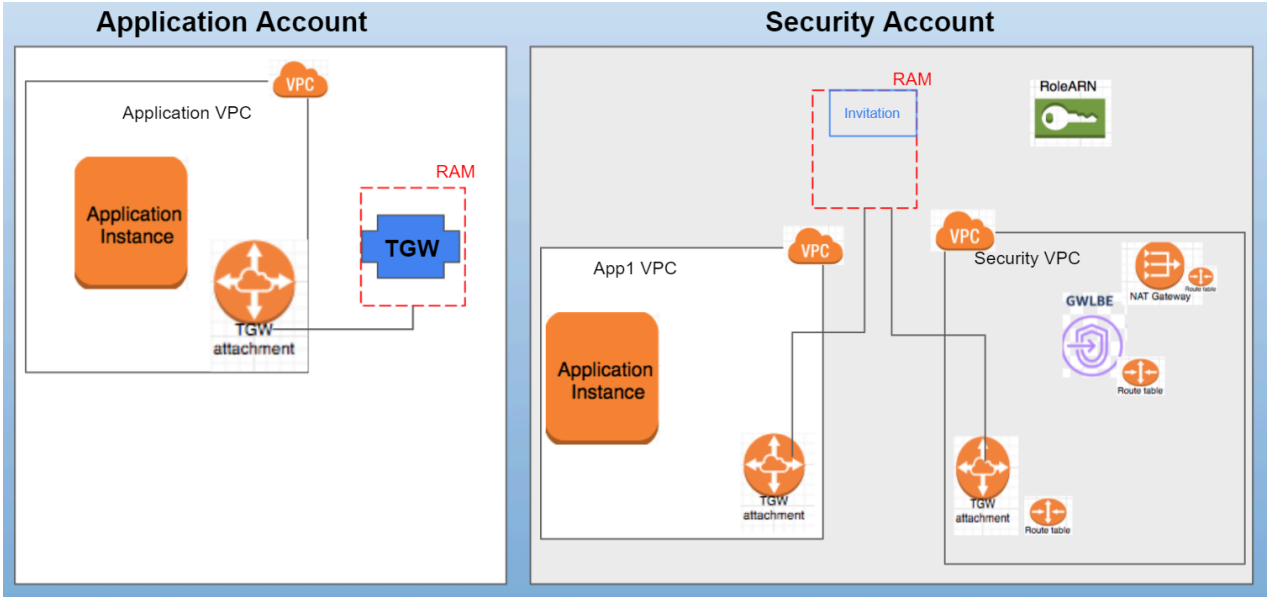

 ©2024 Palo Alto Networks, Inc.

VM-Series Integration with an AWS Gateway Load Balancer

The AWS Gateway Load Balancer (GWLB) is an AWS managed service that allows you to deploy a stack of VM-Series firewalls and operate in a horizontally scalable and fault-tolerant manner. You can then expose the AWS GWLB with the stack of firewalls as a VPC endpoint service for traffic inspection and threat prevention. By creating Gateway Load Balancer endpoints (GWLBE) for the VPC endpoint service, you can easily insert an auto-scaling VM-Series firewall stack in the outbound, east-west, and inbound traffic paths of your applications. VM-Series firewalls and the GWLB use the GENEVE encapsulation to keep your traffic packet headers and payload intact, providing complete visibility of the source's identity to your applications.

The VM-Series firewall supports decryption when deployed behind a GWLB for forward and inbound use cases, including TLS1.2 and TLS1.3 utilizing DHE/ECDHE ciphers.

Deploying the VM-Series firewall behind a GWLB requires you to configure the AWS transit gateway.

The image below describes how the integration of GWLB with VM-Series simplifies your AWS transit gateway(TGW) environments. You attach a centralized security VPC to your transit gateway. The centralized security VPC includes a GWLB to scale and load-balance traffic across the stack of VM-Series firewalls.

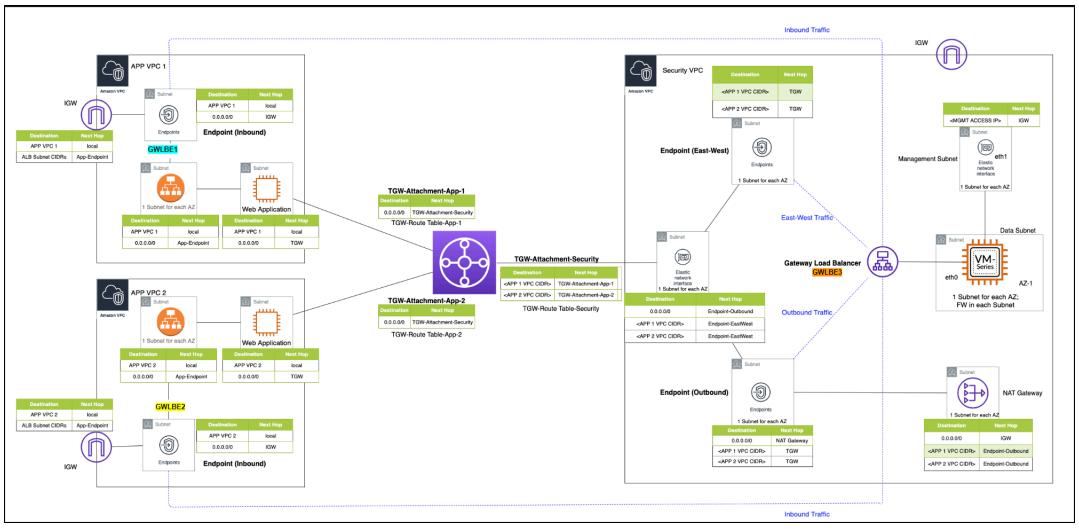

To ensure that the VM-Series firewall can inspect traffic that is routed between VPC attachments, you must enable appliance mode on the transit gateway VPC attachment for the security VPC containing the VM-Series firewall. You can enable appliance mode using the command:

```
modify-transit-gateway-vpc-attachment --transit-gateway-attachment-id
 <value> --options ApplianceModeSupport=enable
```

For more instructions, see enabling appliance mode.

This ensures that bidirectional traffic is routed symmetrically—both request and response traffic are directed to the same Gateway Endpoint in the firewall VPC and the GWLB will maintain

 ©2024 Palo Alto Networks, Inc.

persistence to the same VM-Series firewall for inspection before continuing to the correct destination.

When deployed with a GWLB, you can use the VM-Series firewall to protect:

- Inbound traffic—traffic originating outside the VPC and destined to resources within your application VPC, such as web servers. VM-Series firewalls prevent malware and vulnerabilities from entering the network in traffic allowed by AWS security groups.
- Outbound traffic—traffic originating within the application VPCs and destined to external resources on the Internet. The VM-Series firewalls protect outbound traffic flows by ensuring that workloads in application VPCs connect to permitted services (such as Windows Update) and allowed URL categories and preventing data exfiltration of sensitive information. Additionally, VM-Series security profiles prevent malware and vulnerabilities from entering the network in the return traffic.
- East-West traffic—in a transit gateway environment, East-West traffic refers to Inter-VPC traffic, such as the traffic between source and destination workloads in two different application VPCs. The VM-Series firewalls protect east-west traffic flows against malware propagation.

To protect the inbound traffic to your application VPCs:

1. Create GWLBE endpoints(GWLBE1 and GWLBE2 in the figure above) having separate subnets associated in your spoke VPCs. Ensure that you have separate subnets for GWLB Endpoints, ALB, and Application and Transit Gateway attachment within the application VPC.
2. Add route tables in the application VPC (in addition to the VPC local route) as follows:
    1. Route table with IGW edge association - Add route destined to ALB with target as GWLBE.
    2. Route table with ALB subnet association - Add route destined to 0.0.0.0/0 with target as GWLBE.
    3. Route table with GWLBE subnet association - Add route destined to 0.0.0.0/0 with target as IGW.

With these routes in place, the inbound traffic arriving at VPC IGW is routed towards GWLBE. The GWLBE forwards the traffic to GWLB which in turn sends the traffic to the VM-Series Firewall in the Security VPC for inspection. The firewall sends the request traffic back to the application VPC GWLBE, which then forwards the traffic to the application through ALB. Response traffic to this request is sent by ALB towards the application GWLBE which then sends the traffic to GWLB. The GWLB in turn sends the traffic to the VM-Series firewall. After inspecting the response traffic, the firewall sends the response traffic back to the application GWLBE which in turn sends the traffic to IGW.

To protect the outbound traffic of the application VPCs:

1. Create a GWLBE(GWLBE3 in the figure above) in the centralized firewall VPC. Ensure that you have separate subnets for GWLB Endpoint, Transit Gateway attachment, NAT Gateway within the Security VPC.
2. Create a NAT Gateway in the Security VPC.

 ©2024 Palo Alto Networks, Inc.

3. Add route tables as follows:
   1. Route table with Application subnet association - Add route destined to 0.0.0.0/0 with target as TGW. This is in addition to the VPC local route.
   2. Route tables in Security VPC:
      - Route table with TGW attachment subnet association - In addition to VPC local route, add route destined to 0.0.0.0/0 with target as GWLBE3.
      - Route table with GWLBE subnet association - In addition to VPC local route, add route destined to 0.0.0.0/0 with target as NAT Gateway. Add route destined to Application VPC CIDRs with target as TGW.
      - Route table with NAT Gateway Subnet association - In addition to VPC local route, add route destined to 0.0.0.0/0 with target as IGW. Add route destined to Application VPC CIDRs with target as GWLBE3.
   3. Add Transit Gateway Route tables as follows:
      - Route table with App1-1 VPC TGW-Attachment association - Add route destined to 0.0.0.0/0 with attachment ID as Security VPC TGW attachment.
      - Route table with App2-2 VPC TGW-Attachment association - Add route destined to 0.0.0.0/0 with attachment ID as Security VPC TGW attachment.
      - Route table with Security VPC TGW-Attachment association - (a) Add route destined to App-1 VPC CIDR with attachment ID as Application-1 VPC TGW attachment. (b) Add route destined to App-2 VPC CIDR with attachment ID as Application-2 VPC TGW attachment.

   With this configuration in place, outbound traffic initiated from Application(App1) is sent to TGW and TGW forwards that to the Security VPC subnet. The traffic is then routed to Security GWLBE(GWLBE3) which sends the traffic to VM-Series firewall for inspection through GWLB. The VM-Series firewall sends the traffic back to GWLBE3 after inspection and GWLBE3 forwards the traffic to NAT Gateway which sends the traffic through IGW. Similarly, the response traffic passes through the NAT Gateway to GWLBE3, VM-Series firewall, and TGW after which it is routed back to the application.

   The East-West traffic is also managed with the routes and configuration described in the steps above. When the traffic is sent from App1 to App2, the traffic passes through TGW which routes the traffic to GWLBE3. The GWLBE3 forwards the traffic to the VM-Series firewall through GWLB. The VM-Series firewall sends the packet back to GWLBE3 after inspection. GWLBE3 then forwards the packet to App2 through TGW. The response traffic from App-2 to App-1 will take the reverse path.

   It is recommended to have all subnets in the same AZ to avoid cross-zone traffic charges.

Manual Integration of the VM-Series with a Gateway Load Balancer

Refer to the following topics to manually integrate the VM-Series firewall with an AWS gateway load balancer.

- Enable VM-Series Integration with a Gateway Load Balancer

 ©2024 Palo Alto Networks, Inc.

- Manually Integrate the VM-Series with a Gateway Load Balancer
- (Optional) Associate a VPC Endpoint with a VM-Series Interface
- (Optional) Enable Overlay Routing for the VM-Series on AWS

Enable VM-Series Integration with a Gateway Load Balancer

When integrating the VM-Series firewall with a GWLB, you must first enable the VM-Series firewall to properly process traffic redirected to the firewall by the GWLB endpoints. You can enable this functionality using the VM-Series firewall CLI, through the VM-Series bootstrapping package, or the user-data field in the AWS console.

VM-Series firewall deployment with a GWLB requires:

- PAN-OS 10.0.2 or later
- VM-Series plugin 2.0.2 or later
- Panorama 10.0.2 or later if you using Panorama to manage your firewalls

The table below lists the commands required to enable GWLB traffic inspection with a VPC endpoint. Operation commands can be used in the a bootstrapping init-cfg.txt file or in the user-data field in the AWS console.

Bootstrap Parameter	CLI Command	Description
op-command-modes=mgmt-interface-swap	op-command-modes=mgmt-interface-swap *This command requires the firewall to reboot before taking effect.*	Swaps eth0 and eth1. GWLB by default, sends traffic only to Eth0 of its target instances. By swapping, Eth0 becomes the data interface and eth1 becomes the management interface.
plugin-op-commands=aws-gwlb-inspect:enable	request plugins vm_series aws gwlb inspect enable <yes/no>	Enables the VM-Series firewall to process traffic passing through a GWLB.

Manually Integrate the VM-Series with a Gateway Load Balancer

Complete the following procedure to manually integrate your VM-Series firewall on AWS with a GWLB.

If you associate VPC endpoints to an interface or subinterfaces via user data while bootstrapping and your bootstrap.xml file does not include the interface configuration, you can configure the interfaces after the firewall boots up.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Set up the security VPC. See the AWS documentation for more information about creating your security VPC.

- Create two subnets—one for management and one for data.
- Create two security groups—one for firewall management and one for data.
- The management subnet security groups should allow https and ssh for management access.
- Ensure that the security group(s) in your data VPC allows GENEVE-encapsulated packets (UDP port 6081).
- If your deployment includes a transit gateway and traffic that will move between VPCs, you must enable appliance mode on security VPC attachment.

The target group of the GWLB cannot use HTTP for health checks because the VM-Series firewall does not allow access with an unsecured protocol. Instead, use another protocol such as HTTPS or TCP.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Launch the VM-Series firewall.

1. On the EC2 Dashboard, click **Launch Instance**.
2. Select the VM-Series AMI. To get the AMI, see Obtain the AMI.
3. Launch the VM-Series firewall on an EC2 instance.
   1. Choose the **EC2 instance type** for allocating the resources required for the firewall, and click **Next**. See VM-Series System Requirements, for resource requirements.
   2. Select the security VPC.
   3. Select the data subnet to attach to eth0.
   4. Add another network interface for eth1 to act as the management interface after the interface swap. Swapping interfaces requires a minimum of two ENIs (eth0 and eth1).
      - Expand the Network Interfaces section and click **Add Device** to add another network interface and configure Management subnet to this interface.

        Make sure that your VPC has more than one subnet so that you can add additional ENIs at launch.

        If you launch the firewall with only one ENI:
        - *The interface swap command will cause the firewall to boot into maintenance mode.*
        - *You must reboot the firewall when you add the second ENI.*
      - Expand the Advanced Details section and in the **User data** field enter as text to perform the interface swap during launch.

        `mgmt-interface-swap=enable`

        `plugin-op-commands=aws-gwlb-inspect:enable`

        If you set the target type to the IP address of a specific interface on the VM-Series firewall, you do not need to enable management interface swap.

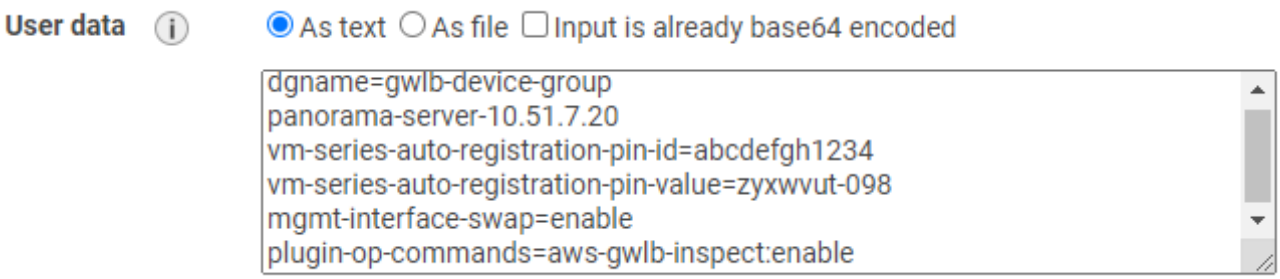

   5. Accept the default **Storage** settings. The firewall uses volume type SSD (gp2).
   6. If prompted, select an appropriate **SSD** option for your setup.
   7. (Optional) . Add one or more tags to create your own metadata to identify and group the VM-Series firewall. For example, add a **Name** tag with a **Value** that helps you remember that the ENI interfaces have been swapped on this VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

8. Select the data **Security Group** for eth0 (data interface). Enable traffic on UDP port 6081.

   If you enable health checks to the firewall, you cannot use HTTP. Instead, use another protocol such as HTTPS or TCP.
9. Select **Review and Launch**. Review that your selections are accurate and click **Launch**.
10. Select an existing key pair or create a new one, and acknowledge the key disclaimer.

    This key pair is required for first time access to the firewall. It is also required to access the firewall in maintenance mode.
11. Download and save the private key to a safe location; the file extension is `.pem`. You cannot regenerate this key, if lost.

    It takes 5-7 minutes to launch the VM-Series firewall. You can view the progress on the EC2 Dashboard.When the process completes, the VM-Series firewall displays on the **Instances** page of the EC2 Dashboard.

STEP 3 | Attach the management security group to eth1 (management interface). Allow ssh and https. See the AWS Documentation for more information.

STEP 4 | Create and assign an Elastic IP address (EIP) to the ENI used for management access (eth1) to the firewall.

1. Select **Elastic IPs** and click **Allocate New Address**.
2. Select **EC2-VPC** and click **Yes, Allocate**.
3. Select the newly allocated EIP and click **Associate Address**.
4. Select the **Network Interface** and the **Private IP address** associated with the management interface and click **Yes, Associate**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Configure a new administrative password for the firewall.

On the VM-Series firewall CLI, you must configure a unique administrative password before you can access the web interface of the firewall. To log in to the CLI, you require the private key that you used to launch the firewall.

1. Use the EIP to SSH into the Command Line Interface (CLI) of the VM-Series firewall. You will need the private key that you used or created above and using the user name **admin** to access the CLI.

   If you are using PuTTY for SSH access, you must convert the .pem format to a .ppk format. See https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/putty.html
2. Enter the following command to log in to the firewall:

   ssh-i <*private_key.pem*> admin@<*public-ip_address*>
3. Configure a new password, using the following command and follow the onscreen prompts:

   configure

   set mgt-config users admin password
4. If you have a BYOL that needs to be activated, set the DNS server IP address so that the firewall can aceess the Palo Alto Networks licensing server. Enter the following command to set the DNS server IP address:

   set deviceconfig system dns-setting servers primary <*ip_address*>
5. Commit your changes with the command:

   commit
6. Terminate the SSH session.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Configure the dataplane network interface as a Layer 3 interface on the firewall.

On the application servers within the VPC, define the dataplane network interface of the firewall as the default gateway.

1. Using a secure connection (https) from your web browser, log in using the EIP address and password you assigned during initial configuration (https://<Elastic_IP address>). You will see a certificate warning; that is okay. Continue to the web page.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default virtual router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone and leave the remaining fields with default values and then click **OK**.
   - On the **IPv4** tab, select **DHCP Client**.

     If using DHCP, select **DHCP Client**; the private IP address that you assigned to the ENI in the AWS management console will be automatically acquired.
   - On the Advanced tab, create a management profile to enable HTTP service as part of management profile creation and allow Health check probes from GWLB.
4. Click **Commit**. Verify that the link state for the interface is up.

STEP 7 | Create security policies to allow/deny traffic.

Because the VM-Series treats traffic as intrazone when integrated with a GWLB, a default intrazone rule allows all traffic. It is a best practice to override the default intrazone rule with a deny action for traffic that does not match any of your other security policy rules.

1. Select **Policies** > **Security** on the web interface of the firewall.
2. Click **Add**, and specify the security zones, applications and logging options that you would like to execute to restrict and audit traffic traversing through the network.

STEP 8 | **Commit** the changes on the firewall.

Associate a VPC Endpoint with a VM-Series Interface

You can associate one or more VPC endpoints with an interface or subinterface of the VM-Series firewall. You can provide consistent policy enforcement by associating all the endpoints in a single VPC to the same subinterface on the firewall. Or, if your deployment has VPCs with overlapping IP address, you can associate endpoints in different VPCs with different subinterfaces for differentiated policy enforcement.

Associating a VPC to an interface or subinterface is not mandatory to integrate the VM-Series firewall with a GWLB.

 ©2024 Palo Alto Networks, Inc.

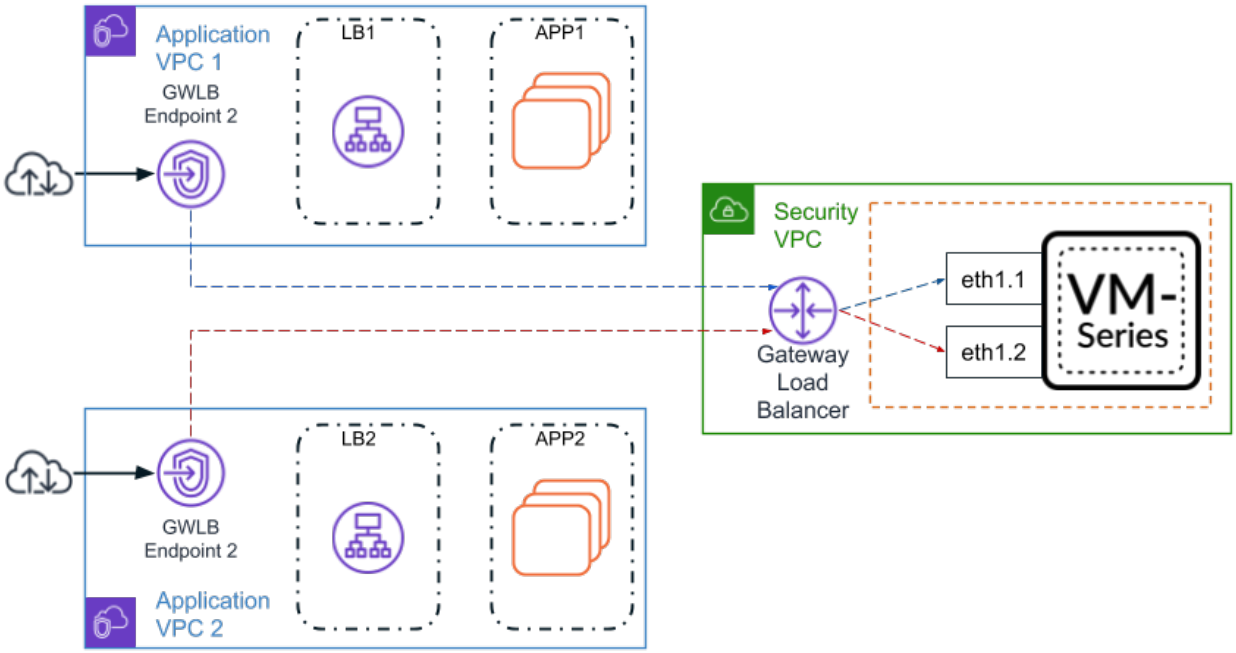

You can configure interfaces and associate a VPC with firewall interfaces using the following methods:

- Include the interface configuration in your `bootstrap.xml` file and the association commands as part of the `init-cfg.txt` file or AWS user-data.
- After deploying the firewall, manually configure your interfaces and use the firewall CLI to associate your VPCs with interfaces.

You can associate multiple VPC endpoints to a single interface on the VM-Series firewall. However, you must associate each VPC endpoint individually. For example, to associate VPC endpoint 1 and VPC endpoint 2 with subinterface ethernet1/1.2, you must execute the association command separately for each VPC endpoint.

The table below describes the commands used to associate a VPC with an interface. You can include the operation command in your `init-cfg.txt` file or in the AWS user-data.

Bootstrap Parameter	CLI Command	Description
plugin-op-commands= aws-gwlb-associate-vpce:<vpce-id>@ethernet<subinterface>	request plugins vm_series aws gwlb associate vpc-endpoint <vpce-id> interface <subinterface>	Associates a VPC endpoint with an interface or subinterface on the firewall. The specified interface is assigned to a security zone.
—	request plugins vm_series aws gwlb disassociate vpc-endpoint <vpce-id> interface <subinterface>	Disassociates a VPC endpoint with an interface or subinterface on the firewall. The specified interface is assigned to a security zone.

 ©2024 Palo Alto Networks, Inc.

Bootstrap Parameter	CLI Command	Description
—	show plugins vm_series aws gwlb	Displays the operating state of the firewall as it relates to your GWLB deployment. It does not display the firewall configuration. For example, if you configure an association to an interface that does not exist, that association is configured but not part of the operating state. Therefore, it is not displayed.

When associating a VPC endpoint using the bootstrapping init-cfg.txt file or AWS user-date, you can list multiple interfaces or subinterfaces together. All the commands must be on a single line in a comma-separated list with no spaces as shown in the following example.

```
plugin-op-commands=aws-gwlb-inspect:enable,aws-gwlb-associate-
vpce:vpce-0913731043b5c0ebc@ethernet1/1.1,aws-gwlb-associate-
vpce:vpce-08207ccb4cb23a1de@ethernet1/1.1,aws-gwlb-associate-
vpce:vpce-07b66cca88821d6e1@ethernet1/1.2,aws-gwlb-associate-
vpce:vpce-0a9a583fdb928492b@ethernet1/1.3
```

If you are using subinterfaces to separate traffic, create a subinterface for each VPC and associate it to a VPC.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Configure the subinterface.

1. Log in to the firewall web interface.
2. Select **Network** > **Interface**.
3. Highlight **ethernet1/1** and click **Add Subinterface**.
4. Enter a numerical suffix (1 to 9,999) to identify the subinterface.
5. Enter a **VLAN Tag** (1 to 4,094) for the subinterface. This field is required but the VLAN is not used.
6. Select **Virtual Router** as default.
7. Select a **Security Zone**.
8. On the **IPv4** tab, set the **Type** to **DHCP Client**.
9. Click **OK**.
10. Repeat this command for each VPC endpoint.

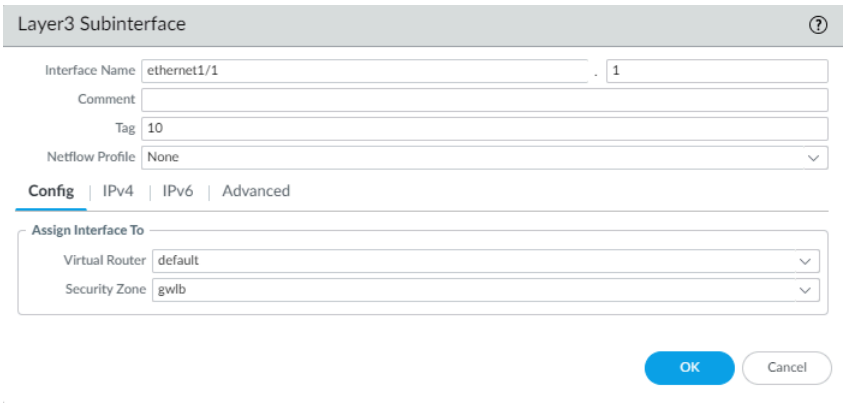

STEP 2 | Associate the interface with a VPC endpoint.

1. Log in to the firewall CLI.
2. Execute the following command:

   request plugins vm_series aws gwlb associate vpc-endpoint <vpce-id> interface <subinterface>

   For example:

   request plugins vm_series aws gwlb associate vpc-endpoint vpce-02c4e6g8ha97h7e39 interface ethernet1/1.4

   You can locate the VPC endpoint ID in the AWS console.

3. Repeat this command for each interface and VPC endpoint association.

STEP 3 | Verify your interface to VPC endpoint associations.

show plugins vm_series aws gwlb

```
GWLB enabled:      True
Overlay Routing:  False
------------------------------------------------------------
VPC endpoint                    Interface
```

 ©2024 Palo Alto Networks, Inc.

```
------------------------------------------------------------
vpce-0aeb1a919bd4ae609           ethernet1/1.1
vpce-0294375bfe413f04a           ethernet1/1.2
```

STEP 4 | If necessary, you can use the following command to disassociate a VPC endpoint from a interface.

```
request plugins vm_series aws gwlb disassociate vpc-endpoint <vpce-id> interface <subinterface>
```

Enable Overlay Routing for the VM-Series on AWS

Overly routing requires PAN-OS 10.0.5 or later.

Using overlay routing in your VM-Series firewall integration the AWS GWLB allows you to use two-zone policy to inspect traffic leaving (egressing) your AWS environment. This allows packets to leave the VM-Series firewall through a different interface than that which they entered through.

When overlay routing is configured, the firewall is able to perform a Layer 3 route lookup a packet's inner header. If the destination is the same as the ingress interface, the packet will be directed as normal. All future packets in the session are treated as vwire; as if overlay routing was not enabled. If the packet is going to an outbound destination, the firewall decapsulates the packet and forwards the packet to the IGW or NAT gateway. When the packet returns, the firewall reapplies the encapsulation.

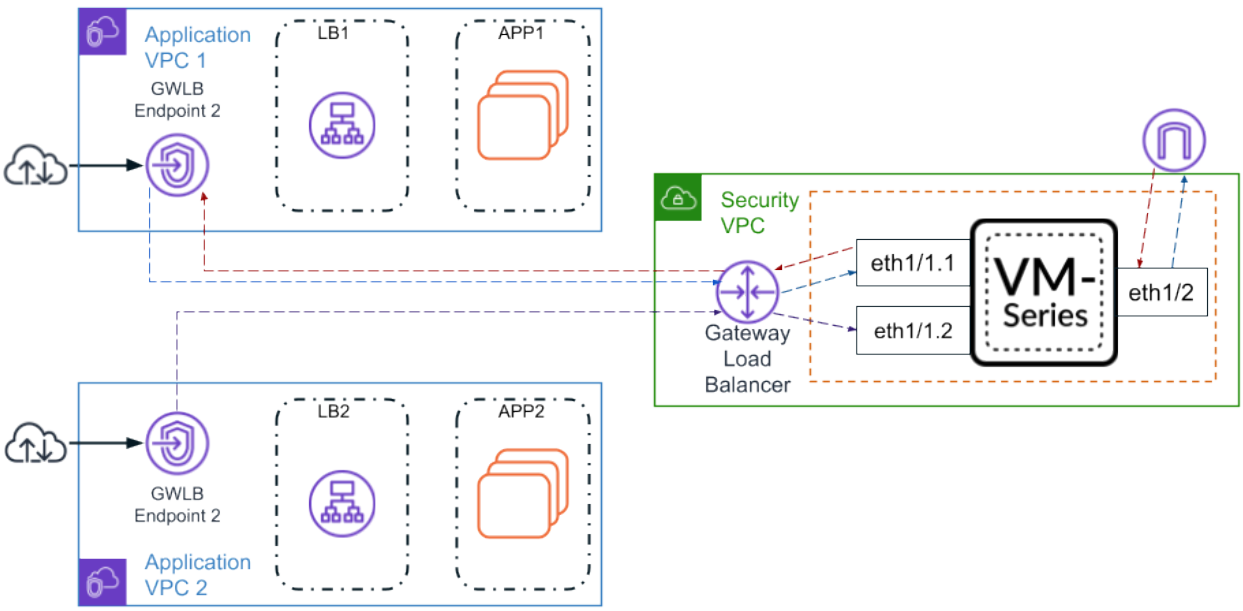

Use the following procedure to enable overlay routing.

STEP 1 | Before you begin, ensure that you create different subnets for the trust and untrust interfaces.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Manually Integrate the VM-Series with a Gateway Load Balancer.

STEP 3 | (Optional) Associate a VPC Endpoint with a VM-Series Interface.

STEP 4 | Use overlay routing CLI command. This CLI command is not required if you included the overlay routing op-command in the AWS user-data or the init-cfg.txt bootstrap file.

1. Log in to the firewall command line interface.
2. Execute the following command.

```
request plugins vm_series aws gwlb overlay-routing enable yes
```

STEP 5 | Log in to the firewall web interface.

STEP 6 | Disable **Automatically create default route pointing to default gateway provided by server** on the trust (ingress) interface.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Click on your trust interface and then the IPv4 tab.
3. Uncheck **Automatically create default route pointing to default gateway provided by server**.
4. Click OK.

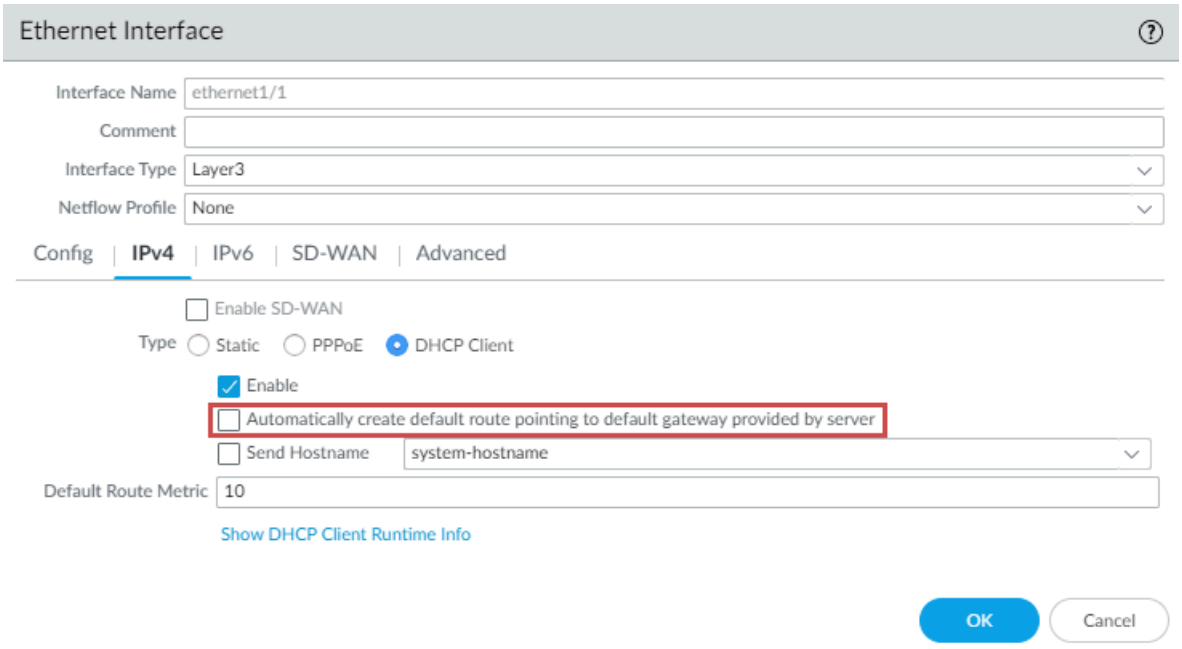

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Configure interface Ethernet 1/2.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Select the **Interface Type**—**Layer 3**.
3. On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. This zone will act as your untrust zone and directing outbound traffic out of your security VPC. Define the new zone, such as VM-Series-untrust, and then click **OK**.
4. On the **IPv4** tab, select **DHCP Client**.
5. Select **Automatically create default route pointing to default gateway provided by server**.
6. Click **OK**.

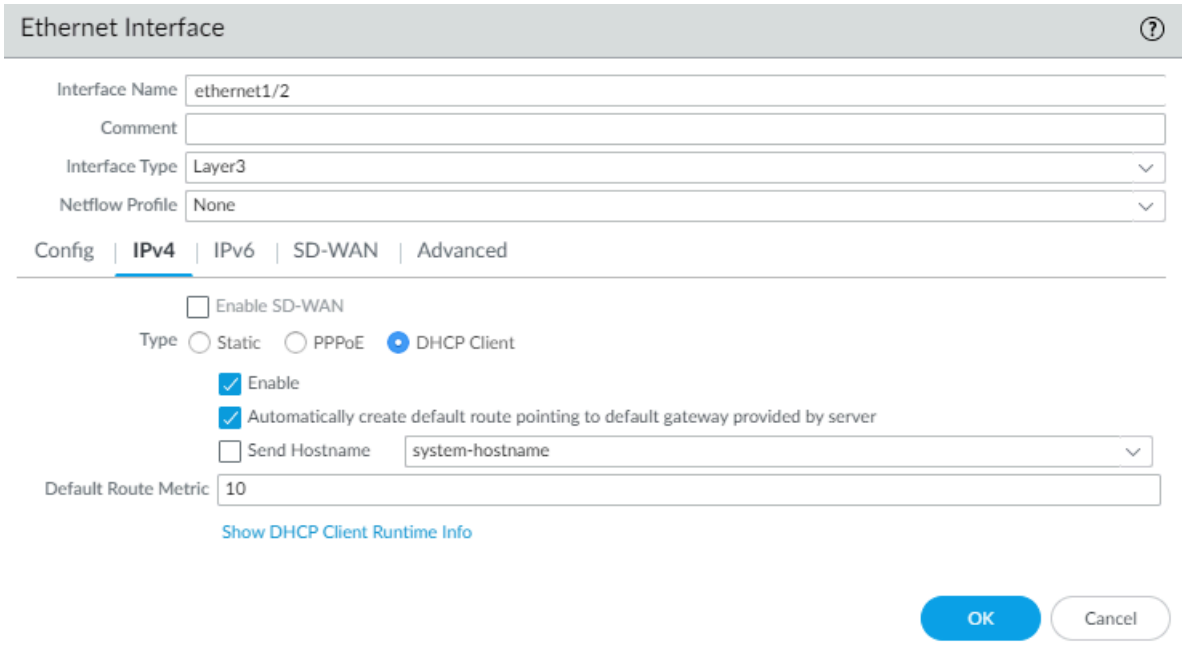

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Configure a virtual router.

1. Select **Network** > **Virtual Routers** > **Add**.
2. Enter a descriptive **Name** for the virtual router.
3. Under **Interfaces**, **Add** Ethernet1/1, any subinterfaces under Ethernet1/1, and Ethernet1/2.
4. Click **Static Routes** > **Add**.
    1. Enter a descriptive name for the static route.
    2. As the **Destination**, enter the private IP address of the application VPC subnet.
    3. Select the trust (ingress) interface from the **Interface** drop-down.
    4. For **Next Hop**, select IP Address and enter the IP address of the gateway of the trust interface. You can find the gateway IP address on **Network** > **Interfaces** > **Ethernet** > **Dynamic-DHCP Client**.

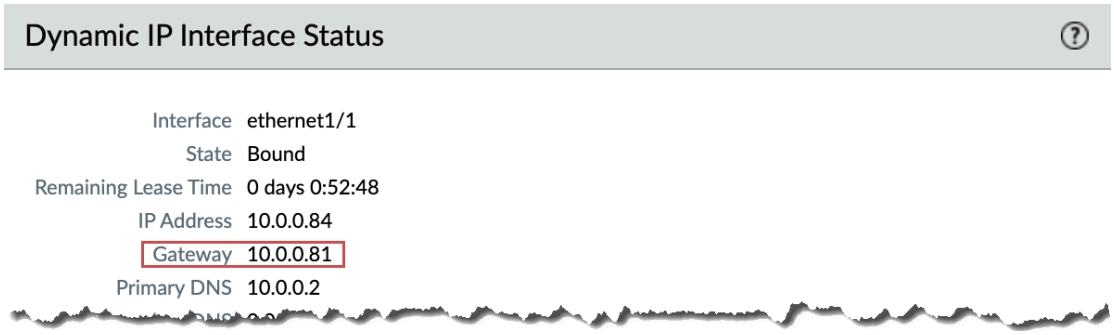

    5. Click **OK**.

 ©2024 Palo Alto Networks, Inc.

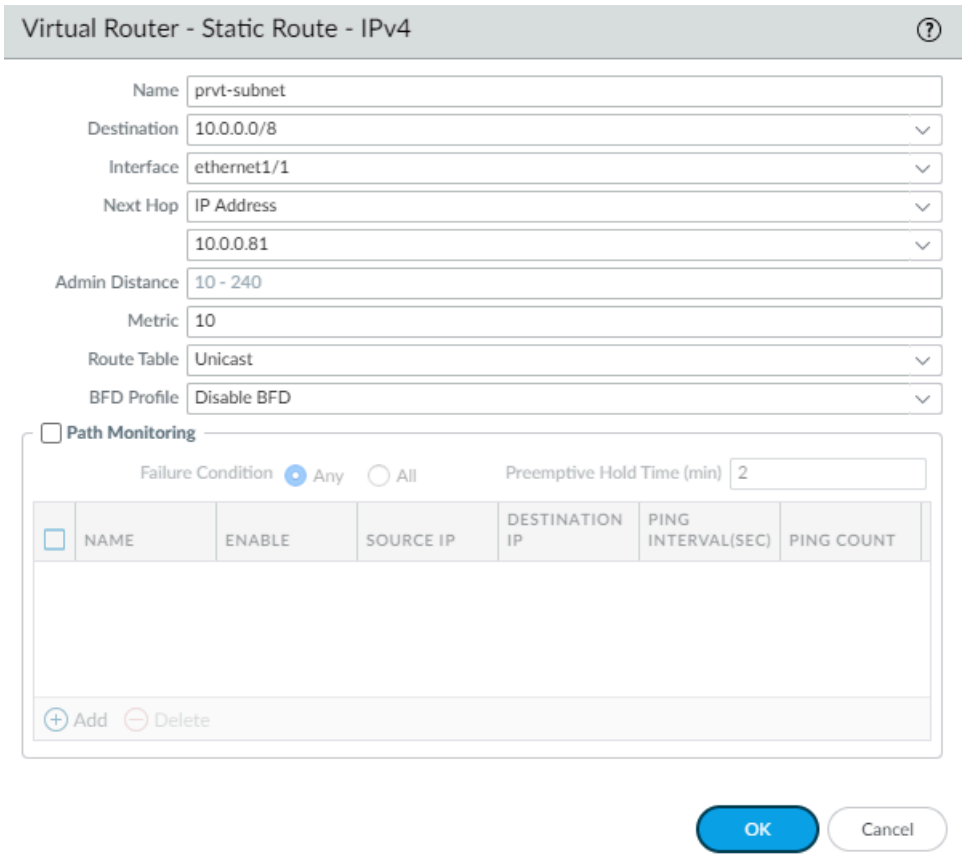

5. Ensure that the static routes can reach all application VPC in your deployment. You can either make a few large aggregated routes (covering all RFC1918) or application VPC specific routes. If you use subinterfaces, you do not need to route back to the sub-interface. The egress check looks only for the matching interface instead of the matching subinterface.
6. Click **OK**.

STEP 9 | Create a NAT policy for traffic egressing Ethernet1/2.

1. Select **Policies** > **NAT** > **Add**.
2. Enter a descriptive **Name** for the NAT policy rule.
3. Select **ipv4** from the **NAT Type** drop-down.
4. On the **Original Packet** tab, set the **Source Zone** to any and the **Destination Zone** to your untrust (egress) zone.
5. On the **Translated Packet** tab, set the following parameters.
   - Translation Type: Dynamic IP and Port
   - Address Type: Interface Address
   - Interface: Select your untrust (egress) port from the drop-down.
   - IP Address: None
6. Click **OK**.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | **Commit** your changes.

VM-Series Auto Scaling Group with AWS Gateway Load Balancer

The Palo Alto Networks auto scaling template for AWS help you integrate and configure the VM-Series firewall with a GWLB to protect applications deployed in AWS. The template leverage AWS scalability features to independently and automatically scale VM-Series firewalls deployed in AWS to meet surges in application workload resource demand.

*These templates are **community supported**.*

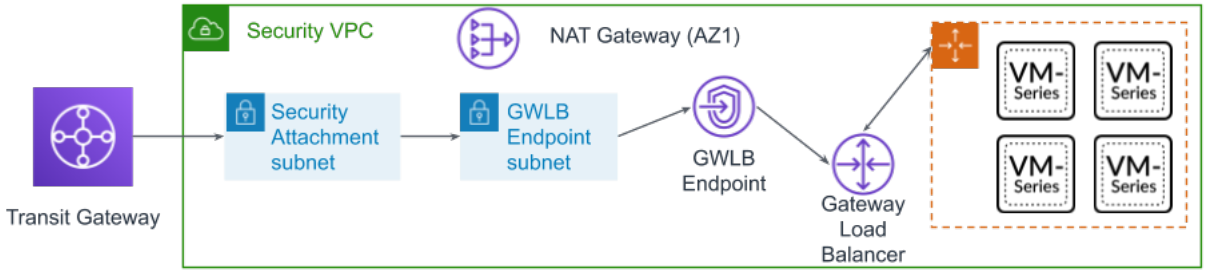

This solution provides a security VPC template and an application template. The security VPC template deploys the VM-Series firewall auto scaling group, a GWLB, a GWLBE, GWLBE subnet, security attachment subnet, and a NAT gateway for each availability zone. Download the CloudFormation templates from the Palo Alto Networks GitHub Repository.

The VM-Series Auto Scaling template for integration with an AWS GWLB includes the following building blocks:

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

Building Block	Description
PAN Components	• Panorama running 10.0.2 or later • PAN-OS 10.0.2 or later • VM-Series plugin 2.0.2 or later installed on Panorama
Firewall template (Community supported template)	Based on the number of availability zones (AZs) you choose, the `firewall-new-vpc-v3.0.template` deploys the following: *The template supports a maximum of four AZs.* • Subnets for Lambda management, transit gateway attachments, GWLB endpoints, and NAT gateways, as well as trust subnets. • Routes tables for each subnet • Transit gateway attachments and route tables

 ©2024 Palo Alto Networks, Inc.

Building Block	Description
	• NAT and internet gateways • An auto scaling group with one VM-Series firewall per AZ. • One GWLB and a GWLB endpoint in each AZ. The VPC CIDR for the firewall template should be larger than /23. Due to the many variations in a production environment that includes but is not limited to a specific number components, such as subnets, availability zones, route tables, and security groups. You must deploy the `firewall-new-vpc-v3.0.template` in a new VPC. *VM-Series Auto Scaling template for AWS does not deploy a transit gateway or Panorama. You must deploy a transit gateway and Panorama before launching firewall-new-vpc-v3.0.template.*
Application template (Community supported template)	Based on the number of availability zones (AZs) you choose, the `panw-aws-app-v3.0.template` deploys the following: *The template supports a maximum of four AZs.* • Subnets for Lambda, transit gateway attachments, GWLB endpoints, application load balancers. • Routes tables for each subnet, as well as an inbound route table associated with the internet gateway to direct inbound traffic to the GWLB endpoint. • One application load balancer • One internet gateway • An auto scaling group with one Ubuntu instance per AZ. The VPC CIDR for the application template should be larger than /23. *The application template is intended to be used as an example for validating the security template.*
Lambda functions	AWS Lambda provides robust, event-driven automation without the need for complex orchestration software. In addition to deploying the components described in the rows above, the `firewall-new-vpc-v3.0.template` performs the following functions: • Adds or removes an interface (ENI) when a firewall is launched or terminated. • Deletes all the associated resources when you delete a stack or terminate an instance. • Removes a firewall as a Panorama managed device when there is a scale-in event.

 ©2024 Palo Alto Networks, Inc.

Building Block	Description
	• Deactivates the license when a scale-in event results in a firewall termination. • Monitors the transit gateway periodically for new attachments or detachments, and updates the route tables accordingly in the security VPC.
Bootstrap files The bootstrap.xml file provided in the GitHub repository is provided for testing and evaluation only. For a production deployment, you must modify the sample credentials in the bootstrap.xml prior to launch.	This solution requires the init-cfg.txt file and the bootstrap.xml file so that the VM-Series firewall has the basic configuration for handling traffic. • The `init-cfg.txt` file includes the mgmt-interface-swap operational command to enable the firewall to receive dataplane traffic on its primary interface (eth0). This auto-scaling solution requires the swapping of the dataplane and management interfaces to enable the GWLB to forward web traffic to the auto-scaling tier of VM-Series firewalls. • The `bootstrap.xml` file enables basic connectivity for the firewall network interfaces and allows the firewall to connect to the AWS CloudWatch namespace that matches the stack name you enter when you launch the template.

If you need to delete these templates from AWS, always delete the application template first. Attempting to delete the firewall template causes the deletion to fail.

- Before Launching the Templates
- Launch the Firewall Template
- Launch the Application Template

Before Launching the Templates

Before you launch the templates to integrate a VM-Series firewall auto scaling group with an AWS GWLB, you must complete the following procedure.

STEP 1 | Ensure that you have the following before you begin.

- Obtain the auth code for a bundle that supports the number of firewalls that might be required for your deployment. You must save this auth code in a text file named

 ©2024 Palo Alto Networks, Inc.

`authcodes` (no extensions), and put the `authcodes` file in the `/license` folder of the bootstrap package.

- Download the files required to launch the VM-Series Gateway Load Balancer template from the GitHub repository.
- Create a Transit Gateway. This transit gateway connects your security and application VPCs.
  - Take note of the transit gateway ID; you will need it later when deploying the template.
  - You must add a 0.0.0.0/0 route to the application attachment route table pointing to the security attachment to protect east-west and outbound traffic.
  - Ensure that **Default route table association** and **Default route table propagation** are disabled.
- The recommended VPC CIDR for the firewall and application templates should be larger than /23.

The target group of the gateway GWLB cannot use HTTP for health checks because the VM-Series firewall does not allow access with an unsecured protocol. Instead use HTTPS or TCP.

STEP 2 | Deploy Panorama running 10.0.2 and configure the following.

Panorama must allow AWS public IP addresses. The VM-Series firewall accesses Panorama using the external IP address of the NAT gateway created by the template.

STEP 3 | Download and install the VM-Series plugin on Panorama.

1. Select **Panorama** > **Plugins** and use **Check Now** to look for new plugin packages. The VM-Series plugin name is `vm_series`.
2. Consult the plugin release notes to determine which version provides upgrades useful to you.
3. Select a version of the plugin and select **Download** in the Action column.
4. Click **Install** in the Action column. Panorama alerts you when the installation is complete.
5. To view the plugin, select **Device** > **VM-Series**.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Configure the template.

1. Log in to the Panorama web interface.
2. Select **Panorama** > **Templates** and click **Add**.
   1. Enter a descriptive **Name**.
   2. Click **OK**.
3. Configure the virtual router.
   1. Select **Network** > **Virtual Routers**.
   2. Ensure that you have selected the template you create above from the **Template** drop-down.
   3. Click **Add**.
   4. Name the virtual router using the following format: VR-<tempstackname>.
   5. Enable ECMP on the virtual router.
   6. Click **OK**.
4. Configure the interface and create the zone.
   1. Select **Network** > **Interfaces** and click **Add Interface**.
   2. Select **Slot 1** and then select the Interface name (for example, ethernet 1/1).
   3. Set **Interface Type** to Layer 3.
   4. On the **Config** tab, select **New Zone** from the **Security Zone** drop-down. In the Zone dialog, define a **Name** for new zone, for example Internet, and then click **OK**.
   5. In the **Virtual Router** drop-down, select virtual router your created above.
   6. Select **IPv4** and click **DHCP Client**.
   7. Click **OK**.
5. Create a management profile that allows HTTPS on the interface created above to support Health Checks.
   1. Select **Network** > **Network Profiles** > **Interface Mgmt** and click **Add**.
   2. Select the protocols that the interface permits for management traffic: **Ping**, **Telnet**, **SSH**, **HTTP**, **HTTP OCSP**, **HTTPS**, or **SNMP**.

      *Don't enable **HTTP** or **Telnet** because those protocols transmit in cleartext and therefore aren't secure.*
6. Assign the Interface Management profile to an interface.
   1. Select **Network** > **Interfaces**, select the type of interface (**Ethernet**, **VLAN**, **Loopback**, or **Tunnel**), and select the interface.
   2. Select **Advanced** > **Other info** and select the Interface **Management Profile** you just added.
   3. Click **OK**.
7. Configure the DNS server and FQDN refresh time.
   1. Select **Device** > **Setup** > **Services** and click the Edit icon.
   2. Set the **Primary DNS Server** to 169.254.169.253. This is the AWS DNS address.
   3. Set the **Minimum FQDN Refresh Time** to 60 seconds.

 ©2024 Palo Alto Networks, Inc.

4. Click **OK**.

8. **Commit** your changes. This is required before proceeding to the next step.
9. Create an administrator.
    1. Select **Device** > **Administrators.**
    2. Enter **`pandemo`** as the **Name**.
    3. Set the **Password** to **`demopassword`** and **Confirm**.
    4. Click **OK**.
10. **Commit** your changes.

STEP 5 | Configure a template stack and add the template to the template stack.

1. Select **Panorama** > **Templates** and **Add Stack**.
2. Enter a unique **Name** to identify the stack.
3. Click **Add** and select the template.
4. Click **OK** to save the template stack.

STEP 6 | Create the **Device Group**.

1. Select **Panorama** > **Device Groups**.
2. Click **Add**.
3. Enter a descriptive **Name**.
4. Click **OK**.
5. Add an allow all security pre-rule.
    1. Select **Policies** > **Security** > **Pre Rules** and click **Add**.
    2. Enter a descriptive **Name**.
    3. Under **Source**, **User**, **Destination**, **Application**, and **Service/URL Category**, select any.
    4. Under **Actions**, select **Allow**.
    5. Click **OK**.
6. **Commit** your changes.

STEP 7 | Add the license deactivation API key for the firewall to Panorama.

1. Log in to the Customer Support Portal.
2. Select **Assets** > **API Key Management**.
3. Copy the API key.
4. Use the CLI to install the API key copied in the previous step.

    `request license api-key set key <key>`

 ©2024 Palo Alto Networks, Inc.

STEP 8 | After deploying Panorama, you must open the following ports as described below on the Panorama security group in AWS.

- **Port 443 (HTTPS)**—Upon initial deployment of the firewall template, leave HTTPS open so Lambda can connect to Panorama.

  When you secure port 443 you specify an IP address range from which you will allow connections, as well as the EIPs assigned to the NAT gateways. The number of NAT gateways in your deployment depends on the number of availability zones you configure. To find NAT gateway EIPs in AWS, go to **VPC** > **NAT Gateways**. Note the EIP information for the security group for HTTPS.

  Additionally, to allow Panorama to release the firewall license after stack deletion, you must allow traffic from the CIDR range of the region where you deployed the firewall template. You can find the CIDR for your region at this link.

- **Port 3978**—Port 3978 must be able to receive traffic from any IP address.

Launch the Firewall Template

This workflow describes how to deploy the firewall template.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Modify the `init-cfg.txt` file and upload it to the `/config` folder.

Because you use Panorama to bootstrap the VM-Series firewalls, your `init-cfg.txt` file should be modified as follows. No `bootstrap.xml` file is needed.

Ensure that you use the device group and template names you created above in the init-cfg.txt file.

```
type=dhcp-client
ip-address=
default-gateway=
netmask=
ipv6-address=
ipv6-default-gateway=
hostname=
vm-auth-key=
panorama-server=
panorama-server-2=
tplname=
dgname=
dhcp-send-hostname=yes
dhcp-send-client-id=yes
dhcp-accept-server-hostname=yesdhcp-accept-server-domain=yes
plugin-op-commands=aws-gwlb-inspect:enable
```

Your init-cfg.txt file must include **`plugin-op-commands=aws-gwlb-inspect:enable`**. This is required when integrating the VM-Series firewall with a GWLB.

You must add the device certificate auto-registration PIN to the init-cfg.txt file to automatically install a device certificate when your VM-Series firewall instance is deployed.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Add the license auth code in the `/license` folder of the bootstrap package.

1. Use a text editor to create a new text file named **`authcodes`** (no extension).
2. Add the authcode for your BYOL licenses to this file, and save. The authcode must represent a bundle, and it must support the number of firewalls that might be required for your deployment. If you use individual authcodes instead of a bundle, the firewall only retrieves the license key for the first authcode in the file.

STEP 3 | Upload Lambda code for the firewall template (`panw-aws.zip`) and the Application template (`app.zip`) to an S3 bucket. You can use the same S3 bucket that you use for bootstrapping.

If the Application stack is managed by a different account than the firewall, use the Application account to create another s3 bucket in the same AWS region as the firewall template and copy `app.zip` to that s3 bucket.

STEP 4 | Select the firewall template.

1. In the AWS Management Console, select **CloudFormation** > **Create Stack**.
2. Select **Upload** the latest firewall template from the Git repository, to choose the firewall template to deploy the resources that the template launches. Click Open and Next.
3. Specify the Stack name. The stack name allows you to uniquely identify all the resources that are deployed using this template.

STEP 5 | Enter a descriptive **Name** for your stack. The name must be 28 characters or less.

STEP 6 | Configure the parameters for the VPC.

1. Enter the number of availability zones and select the region from the availability zone drop-down.
2. Look up the AMI ID for the VM-Series firewall and enter it. Make sure that the AMI ID matches the AWS region, PAN-OS version and the BYOL or PAYG licensing option you opted to use. See Get the Amazon Machine Image IDs for more information.
3. Select the EC2 **Key pair** (from the drop-down) for launching the firewall. To log in to the firewalls, you must provide the name of this key pair and the private key associated with it.
4. Select **Yes** if you want to **Enable Debug Log**. Enabling the debug log generates more verbose logs that help with troubleshooting issues with the deployment. These logs are generated using the stack name and are saved in AWS CloudWatch.

By default, the template uses CPU utilization as the scaling parameter for the VM-Series firewalls. Custom PAN-OS metrics are automatically published to the CloudWatch namespace that matches the stack name you specified earlier.

STEP 7 | Specify the name of the Amazon S3 bucket(s).

1. Enter the name of the S3 bucket that contains the bootstrap package.

   If the bootstrap bucket is not set up properly or if you enter the bucket name incorrectly, the bootstrap process fails, and you cannot log in to the firewall. Health checks for the load balancers also fail.
2. Enter the name of the S3 bucket that contains the panw-aws.zip file. As mentioned earlier you can use one S3 bucket for the Bootstrap and Lambda code.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Specify the keys for enabling API access to the firewall and Panorama.

1. Enter the key that the firewall must use to authenticate API calls. The default key is based on the sample file and you should only use it for testing and evaluation. For a production deployment, you must create a separate PAN-OS login just for the API call and generate an associated key.
2. Enter the API Key to allow AWS Lambda to make API calls to Panorama. For a production deployment, you should create a separate login just for the API call and generate an associated key.

STEP 9 | Add your AWS account number(s). You must provide the account number used to deploy any VPC that is connected to your GWLB. Add these values as a comma-separated list. You can add additional account numbers after deploying the template.

To locate your account number, click your AWS username in the top right of the AWS console and select **My Security Credentials**.

STEP 10 | Enter the transit gateway ID. The transit gateway ID is required to secure east-west and outbound traffic. If you do not enter a transit gateway ID, the template assumes that only inbound traffic should be inspected by firewalls integrated with the GWLB.

STEP 11 | Enter the CIDR for the security VPC.

STEP 12 | Review the template settings and launch the template.

1. Select **I acknowledge that this template might cause AWS CloudFormation to create IAM resources**.
2. Click **Create** to launch the template. The CREATE_IN_PROGRESS event displays.
3. On successful deployment the status updates to CREATE_COMPLETE.

STEP 13 | Verify that the template has launched all required resources.

STEP 14 | Create rules allowing the NAT gateway IP address(es) on the security group where your Panorama appliance is deployed. This is required to allow your firewalls to connect to

 ©2024 Palo Alto Networks, Inc.

Panorama. You can find the list of NAT gateway IP addresses in the CFT security stack output.

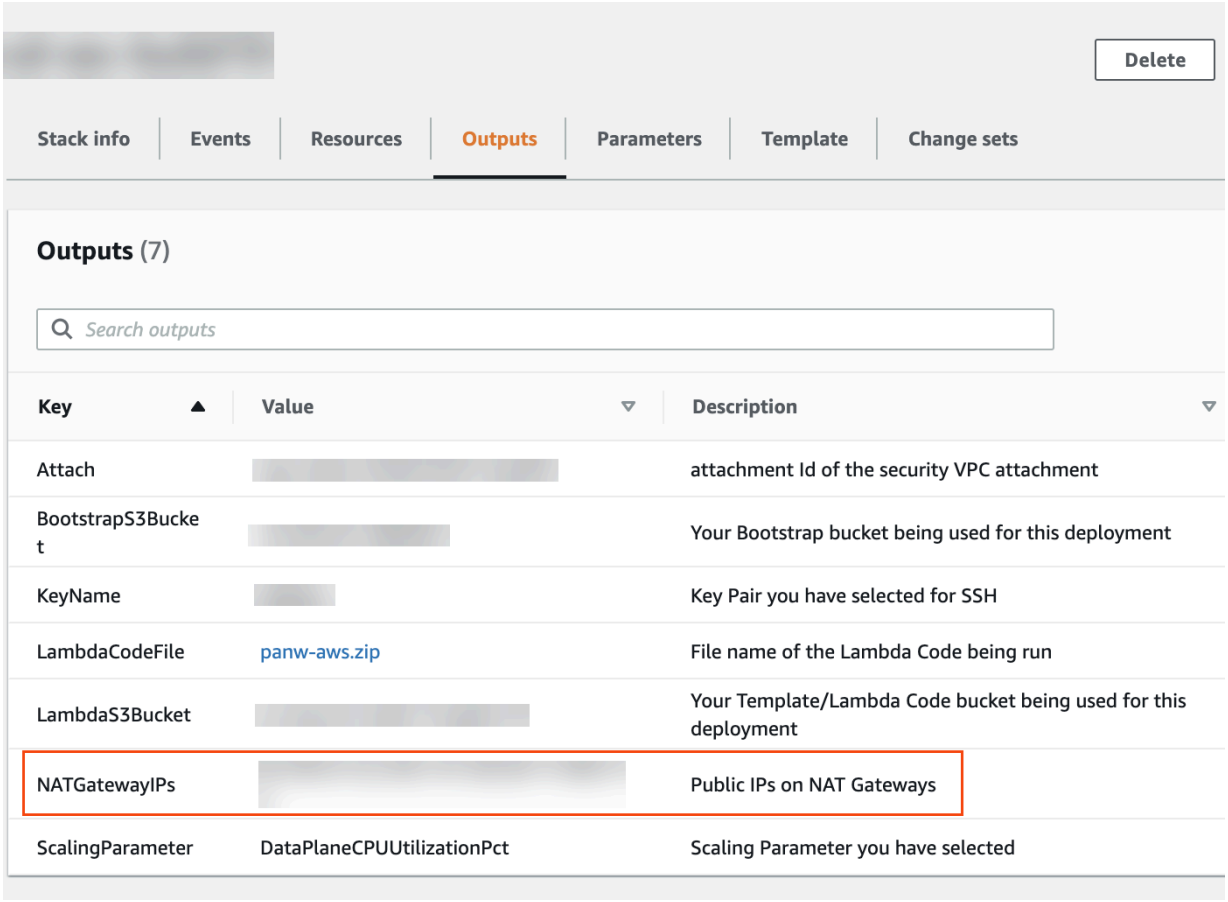

 ©2024 Palo Alto Networks, Inc.

1. Access the AWS VPC console.
2. Select **Security Groups** on the navigation pane.
3. Select the security where Panorama is deployed.
4. Select **Actions** > **Edit Inbound Rules** > **Add rule**.
5. Add rules allowing the NAT gateway IP addresses for Custom TCP Rule for port range 3978.
6. Click **Save rules**.

Type	Protocol	Port Range	Source
HTTP	TCP	80	0.0.0.0/0
Custom TCP Rule	TCP	3978	
Custom TCP Rule	TCP	3978	
Custom TCP Rule	TCP	3978	

Launch the Application Template

Complete the following procedure to launch the application template.

STEP 1 | Create an S3 bucket from which you will launch the application template.

- If this is a cross-account deployment, create a new bucket.
- If there is one account you can create a new bucket or use the S3 bucket you created earlier (you can use one bucket for everything).

STEP 2 | Upload the app.zip file into the S3 bucket.

STEP 3 | Select the application launch template you want you launch.

1. In the AWS Management Console, select **CloudFormation** > **CreateStack**
2. Select Upload a template to Amazon S3, to choose the application template to deploy the resources that the template launches within the same VPC as the firewalls, or to a different VPC. Click **Open** and **Next**.
3. Specify the Stack name. The stack name allows you to uniquely identify all the resources that are deployed using this template.

STEP 4 | Select the Availability Zones (AZ) that your setup will span in Select list of AZ.

STEP 5 | Enter a descriptive **VPC Name**.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Configure the parameters for Lambda.

1. Enter the S3 bucket name where app.zip is stored.
2. Enter the name of the zip file name.

STEP 7 | Select the EC2 instance type for the Ubuntu web server launched by this template.

STEP 8 | Enter your Amazon EC2 key pair.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Enter the name of the service configuration (Service Name) for the GWLB endpoint in the security VPC.

1. Select **DynamoDB** from the **Services** drop-down in the AWS console.
2. Select **Tables** and locate your security VPC table. The table name will be <stack name>-gwlb-<region>. For example—cft-deployment-gwlb-us-east-1.
3. Click the Items tab and copy the Service Name.
4. Paste the Service Name into the application template configuration parameters.

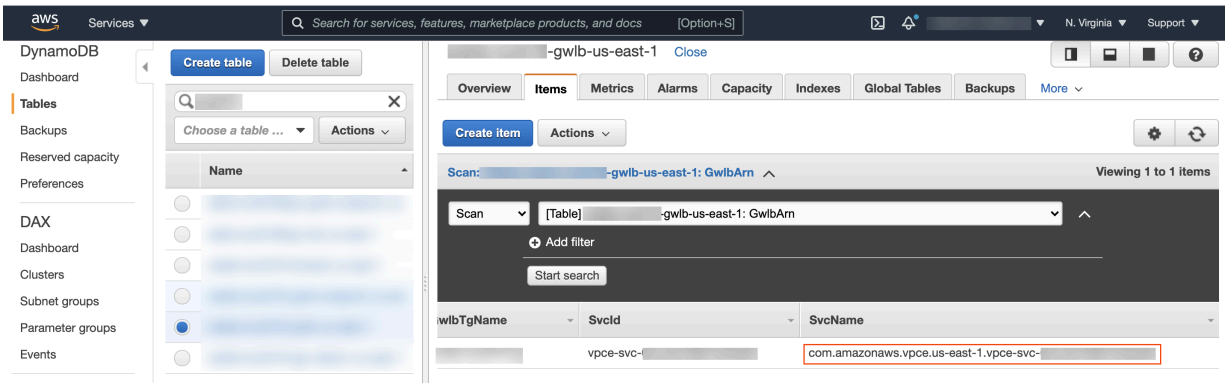

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Enter the transit gateway ID. This is the same transit gateway you created before deploying the firewall template.

STEP 11 | Review the template settings and launch the template.

STEP 12 | After the application has been deployed, you must add a route to the transit gateway route table to enable east-west and outbound traffic inspection.

1. Log in to the AWS VPC console.
2. Select **Transit Gateway Route Tables** and choose your transit gateway route table. This route table is created by the template and is called **<app-stack-name>-<region>-PANWAppAttRt**.
3. Select **Routes** and click **Create static route**.
4. Enter 0.0.0.0/0 in the **CIDR** field.
5. From the **Choose attachment** drop-down, select the VM-Series firewall VPC attachment.
6. Click **Create static route**.

STEP 13 | (Optional) Create a bastion host (also called a jump box) to access the web server created by the application template.

1. Create a public-facing subnet in your application VPC.
2. Add a route to this subnet from your IP address to the internet gateway.
3. Create a new EC2 instance in the public subnet with a public IP address.
4. Create a security group for this EC2 instance that allows SSH from your IP address.

 ©2024 Palo Alto Networks, Inc.

High Availability for VM-Series Firewall on AWS

The VM-Series firewall on AWS supports active/passive HA only; if it is deployed with Amazon Elastic Load Balancing (ELB), it does not support HA (in this case ELB provides the failover capabilities).

- Overview of HA on AWS
- IAM Roles for HA
- HA Links
- Heartbeat Polling and Hello Messages
- Device Priority and Preemption
- HA Timers
- Configure Active/Passive HA on AWS Using a Secondary IP
- Configure Active/Passive HA on AWS Using Interface Move
- Migrate Active/Passive HA on AWS

Overview of HA on AWS

To ensure redundancy, you can deploy the VM-Series firewalls on AWS in an active/passive high availability (HA) configuration. The active peer continuously synchronizes its configuration and session information with the identically configured passive peer. A heartbeat connection between the two devices ensures failover if the active device goes down. There are two options for deploying the VM-Series firewall on AWS in HA—Secondary IP move and Dataplane Interface (ENI) move.

To ensure that all traffic to your internet-facing applications passes through the firewall, you have two options. You can either configure the application's public IP address on the Untrust interface (E1/2 in the illustration above) of the VM-Series firewall, or you can configure AWS ingress routing. The AWS ingress routing capability allows you to associate route tables with the AWS Internet gateway and add route rules to redirect the application traffic through the VM-Series firewall. This redirection ensures that all internet traffic passes through the firewall without having to reconfigure the application endpoints.

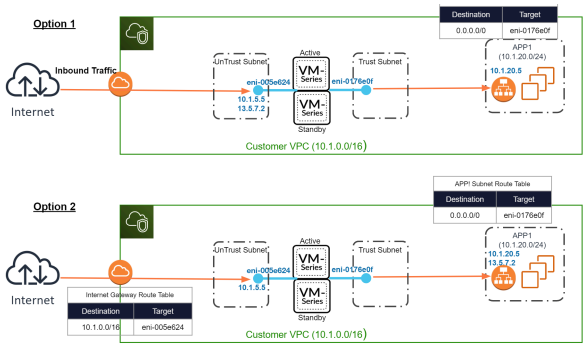

Secondary IP Move

When the active peer goes down, the passive peer detects this failure and becomes active. Additionally, it triggers API calls to the AWS infrastructure to move the configured secondary IP

 ©2024 Palo Alto Networks, Inc.

addresses from the dataplane interfaces of the failed peer to itself. Additionally, AWS updates the route tables to ensure that traffic is directed to the active firewall instance. These two operations ensure that inbound and outbound traffic sessions are restored after failover. This option allows you to take advantage of DPDK to improve the performance of your VM-Series firewall instances and provides better failover time than interface-move HA, while supporting all the features provided by interface-move.

Secondary IP Move HA requires VM-Series plugin 2.0.1 or later.

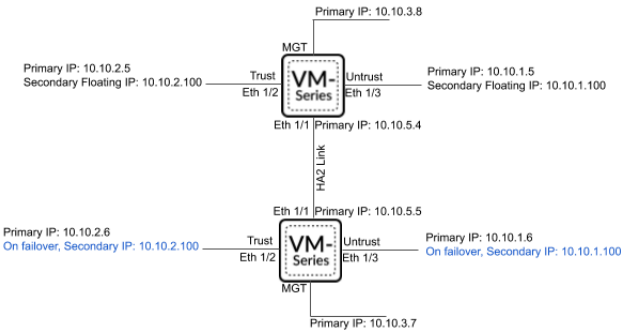

Dataplane Interface Move

When the active peer goes down, the passive peer detects the failure and becomes active. Additionally, it triggers API calls to the AWS infrastructure to move all the dataplane interfaces (ENIs) from the failed peer to itself.

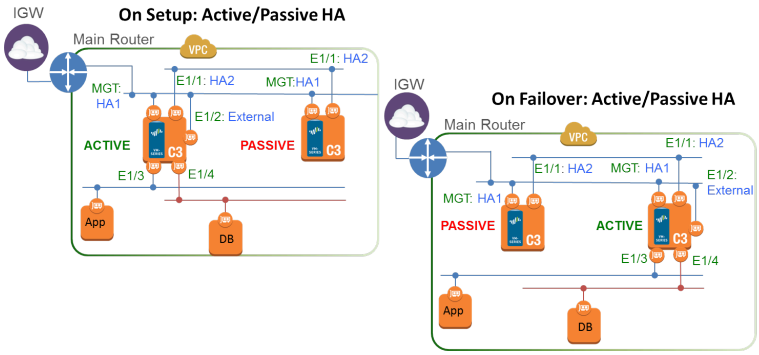

IAM Roles for HA

AWS requires that all API requests must be cryptographically signed using credentials issued by them. In order to enable API permissions for the VM-Series firewalls that will be deployed as an HA pair, you must create a policy and attach that policy to a role in the AWS Identity and Access Management (IAM) service. The role must be attached to the VM-Series firewalls at launch. The policy gives the IAM role permissions for initiating API actions required to move interfaces or secondary IP addresses from the active peer to the passive peer when failover is triggered.

For detailed instructions on creating policy, refer to the AWS documentation on Creating Customer Managed Polices. For detailed instructions on creating an IAM role, defining which accounts or AWS services can assume the role, defining which API actions and resources the

 ©2024 Palo Alto Networks, Inc.

application can use upon assuming the role, refer to the AWS documentation on IAM Roles for Amazon EC2.

The IAM policy, which is configured in the AWS console, must have permissions for the following actions and resources (at a minimum):

The following IAM actions, permissions, and resources are required to enable HA. To enable AWS Cloudwatch monitoring, see Enable CloudWatch Monitoring on the VM-Series Firewall

for the required IAM action.

IAM Action, Permission, or Resource	Description	Interface Move	Secondary IP Move
AttachNetworkInterface	For permission to attach an ENI to an instance.	✓	✓
DescribeNetworkInterfaces	For fetching the ENI parameters in order to attach an interface to the instance.	✓	✓
DetachNetworkInterface	For permission to detach the ENI from the EC2 instance.	✓	✓
DescribeInstances	For permission to obtain information on the EC2 instances in the VPC.	✓	✓
AssociateAddress	For permissions to move public IP addresses associated with the primary IP addresses from the passive to active interfaces.		✓
AssignPrivateIpAddresses	For permissions to assign secondary IP addresses and associated public IP addresses to interfaces on the passive peer.		✓
DescribeRouteTables	For permission to retrieve all route tables associated to the VM-Series firewall instances.		✓
ReplaceRoute	For permissions to update the AWS route table entries.		✓
GetPolicyVersion	For permission to retrieve AWS policy version information.		✓

 ©2024 Palo Alto Networks, Inc.

IAM Action, Permission, or Resource	Description	Interface Move	Secondary IP Move
GetPolicy	For permission to retrieve AWS policy information.		✓
ListAttachedRolePolicies	For permission to retrieve the list of all managed policies attached to a specified IAM role.		✓
ListRolePolicies	For permission to retrieve a list of the names of inline policies embedded in a specified IAM role.		✓
GetRolePolicy	For permission to retrieve a specified inline policy embedded in a specified IAM role.		✓
policy	For permission to access the IAM policy Amazon Resource Name (ARN).		✓
role	For permission to access the IAM roles ARN.		✓
route-table	For permission to access the route table Amazon Resource Name (ARN) to update it upon failover.		✓
Wild card (*)	In the ARN field use the * as a wild card.	✓	✓

The following screenshot shows the access management settings for the IAM role described above for secondary-IP HA:

 ©2024 Palo Alto Networks, Inc.

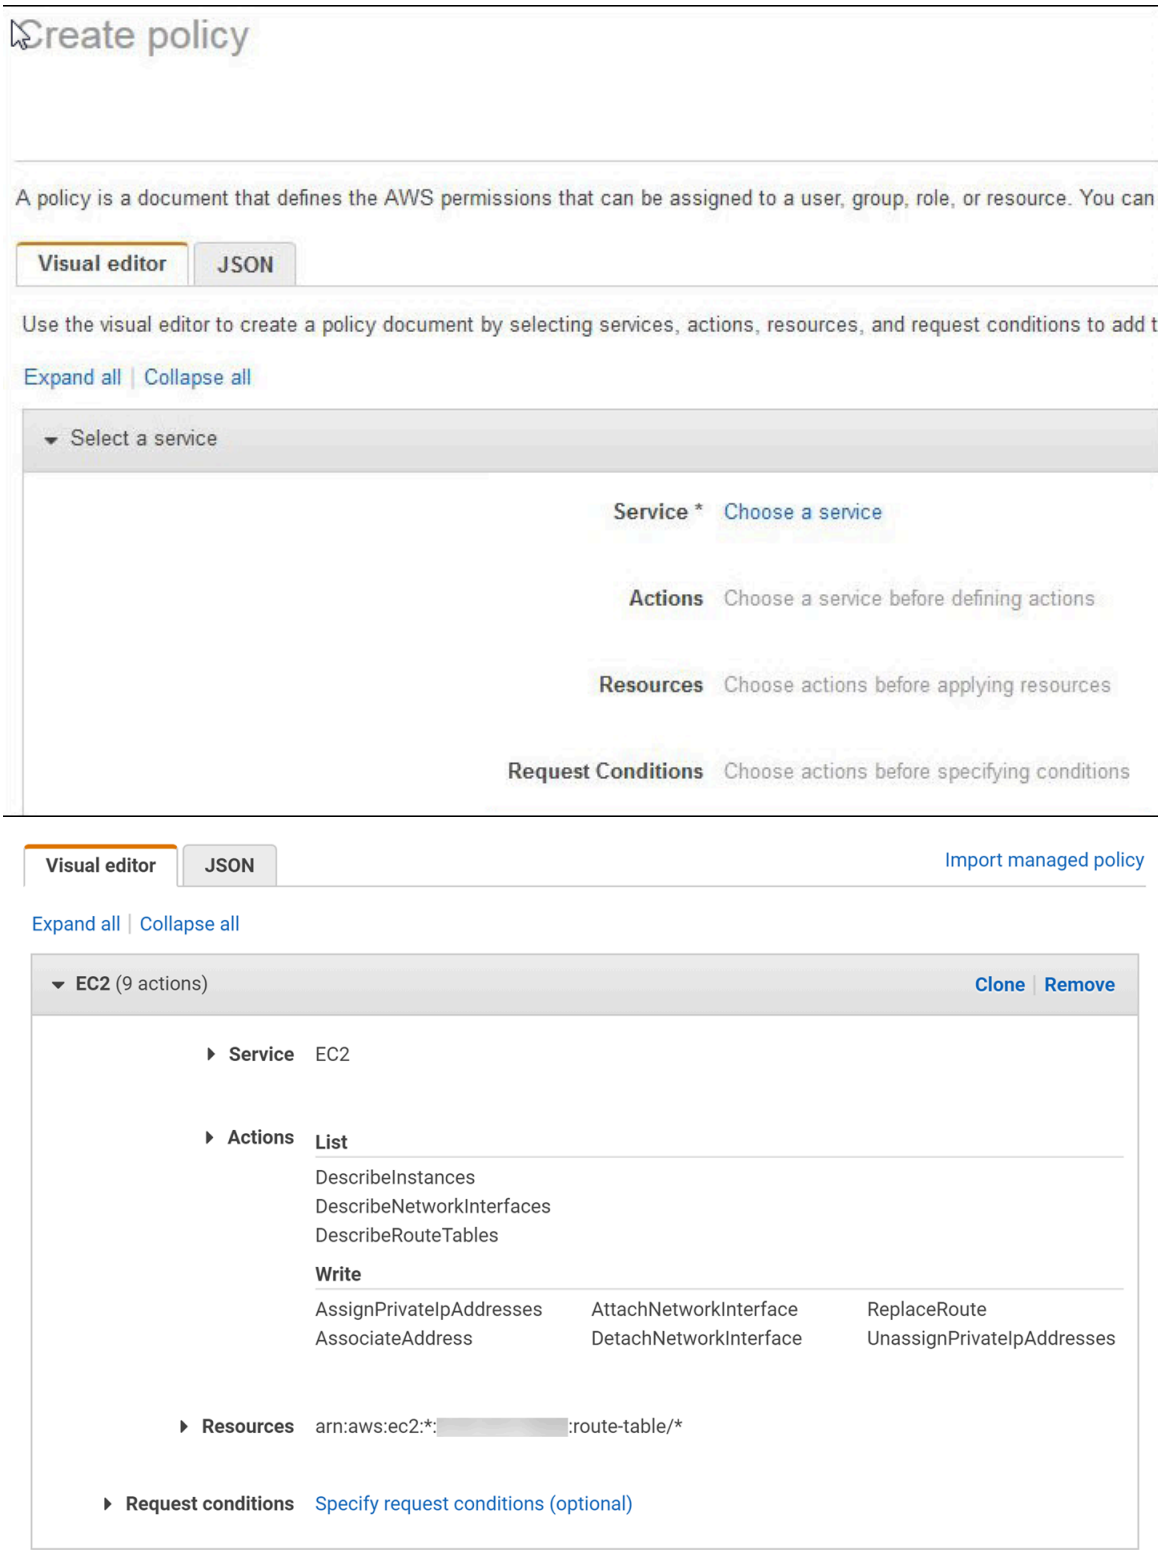

The minimum permissions you need for interface move HA are:

```
{"Version": "2012-10-17","Statement": [{"Sid":
"VisualEditor0","Effect": "Allow","Action":
["ec2:AttachNetworkInterface","ec2:DetachNetworkInterface","ec2:DescribeInstance
"*"}]}
```

 ©2024 Palo Alto Networks, Inc.

The minimum permissions you need for secondary IP move HA are: {"Version": "2012-10-17","Statement": [{"Sid": "VisualEditor0","Effect": "Allow","Action": ["ec2:AttachNetworkInterface","ec2:DetachNetworkInterface","ec2:DescribeInstances","ec2:DescribeNetwork
“ec2:AssignPrivateIpAddresses”,“ec2:AssociateAddress”,“ec2:DescribeRouteTables”],"Resource": "*"}{"Sid": "VisualEditor1","Effect": "Allow","Action": "ec2:ReplaceRoute", "Resource": "arn:aws:ec2:*:*:route-table/*"}]}

HA Links

The devices in an HA pair use HA links to synchronize data and maintain state information. on AWS, the VM-Series firewall uses the following ports:

- **Control Link**—The HA1 link is used to exchange hellos, heartbeats, and HA state information, and management plane sync for routing and User-ID information. This link is also used to synchronize configuration changes on either the active or passive device with its peer.

  The Management port is used for HA1. TCP port 28769 and 28260 for cleartext communication; port 28 for encrypted communication (SSH over TCP).
- **Data Link**—The HA2 link is used to synchronize sessions, forwarding tables, IPSec security associations and ARP tables between devices in an HA pair. Data flow on the HA2 link is always unidirectional (except for the HA2 keep-alive); it flows from the active device to the passive device.

  Ethernet1/1 must be assigned as the HA2 link; this is required to deploy the VM-Series firewall on AWS in HA. The HA data link can be configured to use either IP (protocol number 99) or UDP (port 29281) as the transport.

The VM-Series firewall on AWS does not support backup links for HA1 or HA2.

Heartbeat Polling and Hello Messages

The firewalls use hello message and heartbeats to verify that the peer device is responsive and operational. Hello messages are sent from one peer to the other at the configured *Hello Interval* to verify the state of the device. The heartbeat is an ICMP ping to the HA peer over the control link, and the peer responds to the ping to establish that the devices are connected and responsive. For details on the HA timers that trigger a failover, see HA Timers. (The HA timers for the VM-Series firewall are the same as that of the PA-5200 Series firewalls).

Device Priority and Preemption

The devices in an HA pair can be assigned a *device priority* value to indicate a preference for which device should assume the active role and manage traffic upon failover. If you need to use a specific device in the HA pair for actively securing traffic, you must enable the preemptive behavior on both the firewalls and assign a device priority value for each device. The device with the lower numerical value, and therefore *higher priority*, is designated as active and manages all traffic on the network. The other device is in a passive state, and synchronizes configuration and state information with the active device so that it is ready to transition to an active state should a failure occur.

By default, preemption is disabled on the firewalls and must be enabled on both devices. When enabled, the preemptive behavior allows the firewall with the *higher priority* (lower numerical

 ©2024 Palo Alto Networks, Inc.

value) to resume as active after it recovers from a failure. When preemption occurs, the event is logged in the system logs.

Preemption is not recommended for HA in the VM-Series firewall on AWS.

HA Timers

High availability (HA) timers are used to detect a firewall failure and trigger a failover. To reduce the complexity in configuring HA timers, you can select from three profiles: **Recommended**, **Aggressive**, and **Advanced**. These profiles auto-populate the optimum HA timer values for the specific firewall platform to enable a speedier HA deployment.

Use the **Recommended** profile for typical failover timer settings and the **Aggressive** profile for faster failover timer settings. The **Advanced** profile allows you to customize the timer values to suit your network requirements.

HA Timer on the VM-Series on AWS	Default values for Recommended/Aggressive profiles
Promotion hold time	2000/500 ms
Hello interval	8000/8000 ms
Heartbeat interval	2000/1000 ms
Max number of flaps	3/3
Preemption hold time	1/1 min
Monitor fail hold up time	0/0 ms
Additional master hold up time	500/500 ms

Configure Active/Passive HA on AWS Using a Secondary IP

Complete the following procedure to deploy new VM-Series firewalls as an HA pair with secondary IP addresses.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Before you deploy the VM-Series firewalls for you HA pair, ensure the following:

- Refer to the VPC Planning Worksheet to ensure that your VPC is prepared for the VM-Series firewall.
- Secondary IP Move HA requires VM-Series plugin 2.0.1 or later.
- Deploy both HA peers in the same AWS availability zone.

  Starting with VM-Series plugin 2.0.3, you can deploy the HA peers in different availability zones. Although this type of deployment is not recommended, it is supported.
- Create an IAM role and assign the role to the VM-Series firewalls when you deploy the instances.
- The active and passive firewalls must have at least four interfaces each—a management interface, an HA2 interface, an untrust interface, and a trust interface. Additionally, the trust and untrust interfaces on the active firewall must assigned a secondary IP address.

  The management interface must be used as the HA1 interface.
- Verify that the network and security components are defined suitably.
  - Enable communication to the internet. The default VPC includes an internet gateway, and if you install the VM-Series firewall in the default subnet it has access to the internet.
  - Create subnets. Subnets are segments of the IP address range assigned to the VPC in which you can launch the EC2 instances. The VM-Series firewall must belong to the public subnet so that it can be configured to access the internet.
  - Create a data security group that includes the firewall data interfaces. Additionally, configure the security to allow all traffic (0.0.0.0/0), so security is enforced by the firewalls. This is required to maintain existing sessions during failover.
  - Add routes to the route table for a private subnet to ensure that traffic can be routed across subnets and security groups in the VPC, as applicable.
- If you are bootstrapping the firewall, create the necessary S3 bucket containing the required bootstrap files.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Deploy the VM-Series Firewall on AWS.

1. If your VM-Series firewalls do not have the VM-Series plugin 2.0.1 or later installed, upgrade the plugin before continuing.
2. Configure ethernet 1/1 as the HA2 interface on each HA peer.
    1. Open the Amazon EC2 console.
    2. Select Network Interface and then choose then select your network interface.
    3. Select **Actions** > **Manage IP Addresses.**
    4. Leave the field blank to allow AWS to assign an IP address dynamically or enter an IP address within the subnet range for the VM-Series firewall.
    5. Click **Yes** and **Update**.
    6. Select **Actions** > **Change Source/Dest. Check** and select **Disable**.
    7. Repeat this process on the second (to be passive) HA peer.
3. Add a secondary IP address to your dataplane interfaces on the first (to be active) HA peer.
    1. Select **Network Interface** and then choose then select your network interface.
    2. Select **Actions** > **Manage IP Addresses** > **IPv4 Addresses** > **Assign new IP**.
    3. Leave the field blank to allow AWS to assign an IP address dynamically or enter an IP address within the subnet range for the VM-Series firewall.
    4. Click **Yes** and **Update**.
4. Associate a secondary Elastic (public) IP address with the untrust interface of the active peer.
    1. Select **Elastic IPs** and then choose then select the Elastic IP address to associate.
    2. Select **Actions** > **Associate Elastic IP**.
    3. Under **Resource Type**, select **Network Interface**.
    4. Chose the network interface with which to associate the Elastic IP address.
    5. Click **Associate**.
5. For outbound traffic inspection, add an entry to the subnet route table that sets the next hop as the firewall trust interface.
    1. Select **VPC** > **Route Tables**.
    2. Choose your subnet route table.
    3. Select **Actions** > **Edit routes** > **Add route**.
    4. Enter the **Destination** CIDR Block or IP address.
    5. For **Target**, enter the network interface of the firewall trust interface.
    6. Click **Save routes**.
6. To use AWS Ingress Routing, create a route table and associate the internet gateway to it. Then add an entry with the next hop set as the active firewall untrust interface.
    1. Select **Route Tables** > **Create route table**.
    2. (Optional) Enter a descriptive **Name tag** for your route table.
    3. Click **Create**.

 ©2024 Palo Alto Networks, Inc.

4. Click your route table and select **Actions** > **Edit edge associations**.
5. Select **Internet gateways** and choose your VPC internet gateway.
6. Click **Save**.
7. Click your route table and select **Actions** > **Edit routes**.
8. For the **Target**, select **Network Interface** and choose the untrust interface of the active firewall.
9. Click **Save routes**.

STEP 3 | Configure the interfaces on the firewall. You must configure the HA2 data link and at least two Layer 3 interfaces for your untrust and trust interfaces. Complete this workflow on the first HA peer and then repeat the steps on the second HA peer.

1. Log in to the firewall web interface.
2. Select **Network** > **Interfaces** > **Ethernet** and click on your untrust interface. In this example, the HA2 interface is 1/1, the trust interface is ethernet 1/2, and the untrust interface is ethernet 1/3.
3. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **HA**
4. Click the link for **ethernet 1/2** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example trust-zone, and then click **OK**.
   - On the **IPv4** tab, select **DHCP Client**.
   - Check **Enable**.
   - On the untrust interface, check **Automatically create default route pointing to default gateway provided by server**. This option tells the firewall to create a static route to a default gateway.
   - Repeat these steps for ethernet 1/3.
5. Repeat the above steps on the passive peer.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enable HA.

1. Select **Device** > **High Availability** > **General**.
2. Edit the Setup settings.
3. Enter the private IP address of the passive peer in the **Peer HA1 IP address field**.
4. Click **OK**.

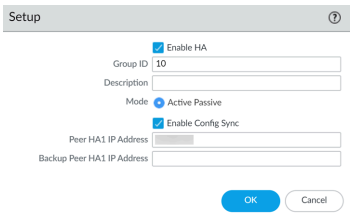

5. Edit the **Election Settings** to specify a particular firewall to be the active peer. Enter a lower numerical **Device Priority** value on the active firewall. If both firewalls have the same Device Priority value, the firewall with the lowest MAC value on the HA1 control becomes the active firewall.

   Enabling preemption is not recommended.

6. Click **OK**.
7. **Commit** your changes.
8. Repeat the above steps on the passive peer.

STEP 5 | Set up the Control Link (HA1) to use the management port.

1. Select **Device** > **High Availability** > **General**, and edit the Control Link (HA1) section.

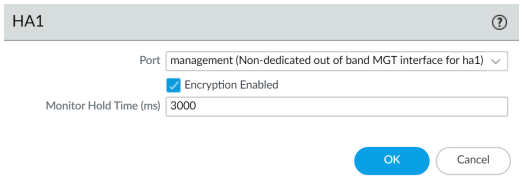

2. (Optional) Select **Encryption Enabled**, for secure HA communication between the peers. To enable encryption, you must export the HA key from a device and import it into the peer device.
   1. Select **Device** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer device can access.
   3. On the peer device, navigate to **Device** > **Certificate Management** > **Certificates**, and select **Import HA key** to browse to the location that you saved the key and import it in to the peer device.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Set up the Data Link (HA2) to use ethernet1/1.

1. Select **Device** > **High Availability** > **General**, edit the Data Link (HA2) section.
2. Select **Port** ethernet1/1.
3. Enter the IP address for ethernet1/1. This IP address must be the same that assigned to the ENI on the EC2 Dashboard.
4. Enter the **Netmask**.
5. Enter a **Gateway** IP address if the HA1 interfaces are on separate subnets.
6. Select **IP** or **UDP** for **Transport**. Use **IP** if you need Layer 3 transport (IP protocol number 99). Use **UDP** if you want the firewall to calculate the checksum on the entire packet rather than just the header, as in the IP option (UDP port 29281).

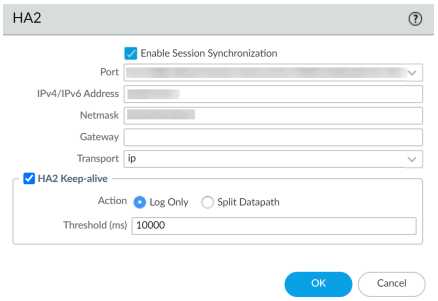

7. (Optional) Modify the **Threshold** for **HA2 Keep-alive** packets. By default, **HA2 Keep-alive** is enabled for monitoring the HA2 data link between the peers. If a failure occurs and this threshold (default is 10000 ms) is exceeded, the defined action will occur. A critical system log message is generated when an HA2 keep-alive failure occurs.

> *You can configure the **HA2 keep-alive** option on both devices, or just one device in the HA pair. If you enable this option on one device, only that device will send the keep-alive messages.*

STEP 7 | After your finish configuring HA on both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls and view the High Availability widget.
2. On the active HA peer, click **Sync to peer**.
3. Confirm that the firewalls are paired and synced.
   - On the passive firewall: the state of the local firewall should display **Passive** and the **Running Config** should show as Synchronized.
   - On the active firewall: the state of the local firewall should display **Active** and the **Running Config** should show as Synchronized.
4. From the firewall command line interface, execute the following commands:
   - To verify failover readiness:

     ```
     show plugins vm_series aws ha state
     ```
   - To show secondary IP mapping:

     ```
     show plugins vm_series aws ha ips
     ```

 ©2024 Palo Alto Networks, Inc.

Configure Active/Passive HA on AWS Using Interface Move

Complete the following procedure to configure active-passive HA using interface move mode.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Make sure that you have followed the prerequisites.

For deploying a pair of VM-Series firewalls in HA in the AWS cloud, you must ensure the following:

- Select the IAM role you created when launching the VM-Series firewall on an EC2 instance; you cannot assign the role to an instance that is already running. See IAM Roles for HA.

  For detailed instructions on creating an IAM role, defining which accounts or AWS services can assume the role, and defining which API actions and resources the application can use upon assuming the role, refer to the AWS documentation.
- DPDK is not supported on the VM-Series firewall on AWS in an interface-move HA deployment. If your have VM-Series plugin 2.0.1 or later on your firewalls, you must disable DPDK.

  Disabling DPDK requires the firewall to reboot. If you are using bootstrapping to deploy the VM-Series firewall, you can avoid rebooting the firewall by disabling DPDK in the initi-cfg.txt file by using **`op-cmd-dpdk-pkt-io=off`**. See Bootstrap the VM-Series Firewall on AWS for more information.
- The active firewall in the HA pair must have at a minimum three ENIs: two dataplane interfaces and one management interface.

  The passive firewall in the HA pair, must have one ENI for management, and one ENI that functions as dataplane interface; you will configure the dataplane interface as an HA2 interface.

  Do not attach additional dataplane interfaces to the passive firewall in the HA pair. On failover, the dataplane interfaces from the previously active firewall are moved —detached and then attached—to the now active (previously passive) firewall.
- The HA peers must be deployed in the same AWS availability zone. While VM-Series HA Across AWS Availability Zones is not a recommended solution, it is supported.

STEP 2 | Launch the VM-Series Firewall on AWS.

STEP 3 | (VM-Series plugin 2.0.1 or later) Disable DPDK on the active and passive firewalls. DPDK is enabled by default and interface-move HA mode does not support DPDK, so you must disable it; enabling Packet MMAP.

1. Log in to the passive firewall CLI.
2. Disable DPDK using the following command. Executing this command restarts the firewall.

   ```
   admin@PA-VM> set system setting dpdk-pkt-io off
   ```

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enable HA.

1. Select **Device** > **High Availability** > **General**, and edit the Setup section.
2. Select **Enable HA**.

STEP 5 | Configure ethernet 1/1 as an HA interface. This interface must be used for HA2 communication.

1. Select **Network** > **Interfaces**.
2. Confirm that the link state is up on ethernet1/1.
3. Click the link for ethernet1/1 and set the **Interface Type** to HA.

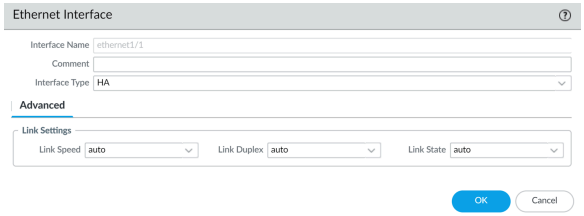

STEP 6 | Set up the Control Link (HA1) to use the management port.

1. Select **Device** > **High Availability** > **General**, and edit the Control Link (HA1) section.

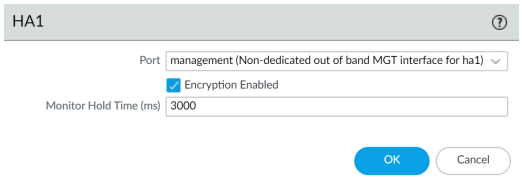

2. (Optional) Select **Encryption Enabled**, for secure HA communication between the peers. To enable encryption, you must export the HA key from a device and import it into the peer device.
   1. Select **Device** > **Certificate Management** > **Certificates**.
   2. Select **Export HA key**. Save the HA key to a network location that the peer device can access.
   3. On the peer device, navigate to **Device** > **Certificate Management** > **Certificates**, and select **Import HA key** to browse to the location that you saved the key and import it in to the peer device.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Set up the Data Link (HA2) to use ethernet1/1.

1. Select **Device** > **High Availability** > **General**, edit the Data Link (HA2) section.
2. Select **Port** ethernet1/1.
3. Enter the IP address for ethernet1/1. This IP address must be the same that assigned to the ENI on the EC2 Dashboard.
4. Enter the **Netmask**.
5. Enter a **Gateway** IP address if the HA1 interfaces are on separate subnets.
6. Select **IP** or **UDP** for **Transport**. Use **IP** if you need Layer 3 transport (IP protocol number 99). Use **UDP** if you want the firewall to calculate the checksum on the entire packet rather than just the header, as in the IP option (UDP port 29281).

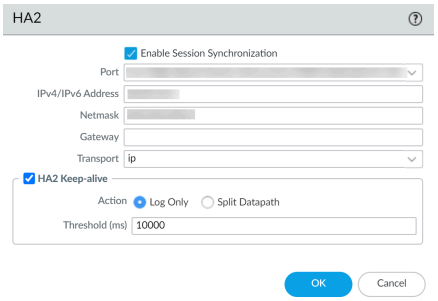

7. (Optional) Modify the **Threshold** for **HA2 Keep-alive** packets. By default, **HA2 Keep-alive** is enabled for monitoring the HA2 data link between the peers. If a failure occurs and this threshold (default is 10000 ms) is exceeded, the defined action will occur. A critical system log message is generated when an HA2 keep-alive failure occurs.

> *You can configure the **HA2 keep-alive** option on both devices, or just one device in the HA pair. If you enable this option on one device, only that device will send the keep-alive messages.*

STEP 8 | Set the device priority and enable preemption.

Use this setting if you want to make sure that a specific device is the preferred active device. For information, see Device Priority and Preemption.

1. Select **Device** > **High Availability** > **General** and edit the Election Settings section.
2. Set the numerical value in **Device Priority**. Make sure to set a lower numerical value on the device that you want to assign a higher priority to.

> *If both firewalls have the same device priority value, the firewall with the lowest MAC address on the HA1 control link will become the active device.*

3. Select **Preemptive**.

   You must enable preemptive on both the active and the passive device.
4. Modify the failover timers. By default, the HA timer profile is set to the **Recommended** profile and is suited for most HA deployments.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | (Optional) Modify the wait time before a failover is triggered.

1. Select **Device** > **High Availability** > **General** and edit the Active/Passive Settings.
2. Modify the **Monitor fail hold up time** to a value between 1-60 minutes; default is 1 minute. This is the time interval during which the firewall will remain active following a link failure. Use this setting to avoid an HA failover triggered by the occasional flapping of neighboring devices.

STEP 10 | Configure the IP address of the HA peer.

1. Select **Device** > **High Availability** > **General**, and edit the Setup section.
2. Enter the IP address of the HA1 port on the peer. This is the IP address assigned to the management interface (ethernet 0/0), which is also the HA1 link on the other firewall.
3. Set the **Group ID** number between 1 and 63. Although this value is not used on the VM-Series firewall on AWS, but cannot leave the field blank.

STEP 11 | Configure the other peer.

Repeat steps 3 to 9 on the HA peer.

STEP 12 | After you finish configuring both devices, verify that the devices are paired in active/passive HA.

1. Access the **Dashboard** on both devices, and view the **High Availability** widget.
2. On the active device, click the **Sync to peer** link.
3. Confirm that the devices are paired and synced, as shown below:
   - On the passive device: The state of the local device should display **passive** and the configuration is **synchronized**.
   - On the active device: The state of the local device should display **active** and the configuration is **synchronized**.

STEP 13 | Verify that failover occurs properly.

1. Verify the HA mode.
   ```
   show plugins vm_series aws ha failover-mode
   ```
2. Verify that the packet IO mode is set to packet MMAP.
   ```
   show system setting dpdk-pkt-io
   ```
3. Shut down the active HA peer.
   1. On the EC2 Dashboard, select **Instances**.
   2. From the list, select the VM-Series firewall and click **Actions** > **Stop**.
4. Check that the passive peer assumes the role of the active peer and that the dataplane interfaces have moved over to the now active HA peer.

Migrate Active/Passive HA on AWS

Both high availability modes are supported, allowing you to migrate between each mode if your deployment requires it. Because interface-move mode does not support DPDK, you must disable it on the VM-Series firewall before completing the migration. Disabling DPDK requires you to reboot the VM-Series firewall, which will impact any traffic sessions on the active firewall.

 ©2024 Palo Alto Networks, Inc.

- Migrate Active/Passive HA on AWS to Secondary IP Mode
- Migrate Active/Passive HA on AWS to Interface Move Mode

Migrate Active/Passive HA on AWS to Secondary IP Mode

Complete the following procedure to migrate your existing VM-Series firewall HA pair from interface-move HA to secondary-IP HA.

Secondary IP Move HA requires VM-Series plugin 2.0.1 or later.

STEP 1 | Upgrade the VM-Series Plugin on the passive HA peer and then the active peer.

STEP 2 | Create secondary IP address for all data interfaces on the active peer.

1. Log in to the AWS EC2 console.
2. Select **Network Interface** and then choose then select your network interface.
3. Select **Actions** > **Manage IP Addresses** > **IPv4 Addresses** > **Assign new IP**.
4. Leave the field blank to allow AWS to assign an IP address dynamically or enter an IP address within the subnet range for the VM-Series firewall.
5. Click **Yes** and **Update**.

STEP 3 | Associate a secondary Elastic (public) IP address with the untrust interface of the active peer.

1. Log in to the AWS EC2 console.
2. Select **Elastic IPs** and then choose then select the Elastic IP address to associate.
3. Select **Actions** > **Associate Elastic IP**.
4. Under **Resource Type**, select **Network Interface**.
5. Chose the network interface with which to associate the Elastic IP address.
6. Click **Associate**.

STEP 4 | Create a route table pointing the subnet containing the trust interface.

1. Select **Route Tables** > **Create route table**.
2. (Optional) Enter a descriptive **Name tag** for your route table.
3. Select your **VPC**.
4. Click **Create**.
5. Select **Subnet Associations** > **Edit subnet associations**.
6. Select the **Associate** checkbox for the subnet containing the trust interface.
7. Click **Save**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Update the IAM roles with additional actions and permissions required to migrate to secondary IP move HA.

IAM Action, Permission, or Resource	Description
AssociateAddress	For permissions to move public IP addresses associated with the primary IP addresses from the passive to active interfaces.
AssignPrivateIpAddresses	For permissions to move secondary IP addresses and associated public IP addresses from the passive to active interfaces.
UnassignPrivateIpAddress	For permissions to unassign secondary IP addresses and associated public IP addresses from interfaces on the active peer.
DescribeRouteTables	For permission to retrieve all route tables associated to the VM-Series firewall instances.
ReplaceRoute	For permission to update the AWS route table entries.
GetPolicyVersion	For permission to retrieve AWS policy version information.
GetPolicy	For permission to retrieve AWS policy information.
ListAttachedRolePolicies	For permission to retrieve the list of all managed policies attached to a specified IAM role.
ListRolePolicies	For permission to retrieve a list of the names of inline policies embedded in a specified IAM role.
GetRolePolicy	For permission to retrieve a specified inline policy embedded in a specified IAM role.
policy	For permission to access the IAM policy Amazon Resource Name (ARN).
role	For permission to access the IAM roles ARN.
route-table	For permission to access the route table ARN.
Wild card (*)	In the ARN field use the * as a wild card.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Create new interfaces (ENIs) on the passive firewall in the same subnet as the active firewall data interfaces.

Do not assign secondary IP addresses to these new interfaces.

1. Open the Amazon EC2 console.
2. Select **Network Interfaces** > **Create Network Interfaces**.
3. Enter a descriptive **Name** for your new interface.
4. Under **Subnet**, select the subnet of the untrust interface of the active firewall.
5. Under **Private IP**, leave the field blank to allow AWS to assign an IP address dynamically or enter an IP address within the subnet range for the untrust interface of the active firewall.
6. Under **Security groups**, select one or more security groups.
7. Select **Yes** and **Create**.
8. Select **Actions** > **Change Source/Dest. Check** and select **Disable**.
9. Repeat these steps for the subnet of the trust interface of the active firewall.

STEP 7 | Attach the new ENIs to the passive firewall instance. You must attach these ENIs to the passive firewall in the correct order because the secondary IP HA method is based on the network interface index assigned by AWS. For example, if eth1/2 on the active firewall is part of subnet A and eth1/3 is part of subnet B, then you must attach the interface that is part of subnet A and the interface that is part of subnet B. In this example, AWS has assigned an index value of 2 to eth1/2 and a value of 3 to eth1/3. This indexing must be maintained for the failover to occur successfully.

1. To attach the ENIs created above, select the untrust interface your created and click **Attach**.
2. Select the Instance ID of the of the passive firewall and click **Attach**.
3. Repeat these steps for the trust interface.

STEP 8 | Log into the passive and set the interfaces to get their IP addresses through DHCP.

1. Log in to the passive VM-Series firewall web interface.
2. Select **Network** > **Interfaces**.
3. Click on the first data interface.
4. Select **IPv4**.
5. Select **DHCP Client**.
6. On the untrust interface only, select **Automatically create default route pointing to default gateway provided by server**.
7. Click OK.
8. Repeat this process for each data interface.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | If you have configure any NAT policies on the VM-Series firewall that reference the private IP addresses of the data interfaces, those policies must be updated to reference the newly assigned secondary IP addresses instead.

1. Access the web interface of the active VM-Series firewall.
2. Select **Policies** > **NAT**.
3. Click on the NAT policy rule to be modified and then **Translated Packet**.
4. Under **Translated Address**, click Add and enter the secondary IP address created in AWS.
5. Delete the primary IP address.
6. Click **OK**.
7. Repeat these steps as necessary.
8. **Commit** your changes.

STEP 10 | Enable secondary IP HA failover mode.

1. Access the VM-Series firewall CLI on the active peer.
2. Execute the following command.

   ```
   request plugins vm_series aws ha failover-mode secondary-ip
   ```
3. Commit your changes.
4. Comfirm your HA mode by executing the following command.

   ```
   show plugins vm_series aws ha failover-mode
   ```
5. Repeat this command on the passive peer.

STEP 11 | After your finish configuring HA on both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls and view the High Availability widget.
2. On the active HA peer, click **Sync to peer**.
3. Confirm that the firewalls are paired and synced.
   - On the passive firewall: the state of the local firewall should display **Passive** and the **Running Config** should show as Synchronized.
   - On the active firewall: the state of the local firewall should display **Active** and the **Running Config** should show as Synchronized.
4. From the firewall command line interface, execute the following commands:
   - To verify failover readiness:

     ```
     show plugins vm_series aws ha state
     ```
   - To show secondary IP mapping :

     ```
     show plugins vm_series aws ha ips
     ```

Migrate Active/Passive HA on AWS to Interface Move Mode

Complete the following procedure to migrate your existing VM-Series firewall HA pair from secondary-IP HA to interface-move HA.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Disable DPDK support on the passive HA peer. Interface-move HA mode does not support DPDK, so you must disable it; enabling Packet MMAP.

1. Log in to the passive firewall CLI.
2. Disable DPDK using the following command. Executing this command restarts the firewall.

   ```
   admin@PA-VM> set system setting dpdk-pkt-io off
   ```

STEP 2 | Disable DPDK support on the active HA peer.

1. Log in to the active firewall CLI.
2. Disable DPDK using the following command. Executing this command restarts the firewall.

   ```
   admin@PA-VM> set system setting dpdk-pkt-io off
   ```

   Restarting the firewall will impact traffic.

STEP 3 | Change the HA mode on the active peer from secondary-IP mode to interface-move mode.

1. Access the VM-Series firewall CLI on the active peer.
2. Execute the following command.

   ```
   request plugins vm_series aws ha failover-mode interface-move
   ```
3. Commit your changes.
4. Comfirm your HA mode by executing the following command.

   ```
   show plugins vm_series aws ha failover-mode
   ```
5. Repeat this command on the passive peer.

STEP 4 | Delete the data interfaces from the passive firewall instance.

1. Log in to the AWS EC2 console.
2. Select **Network Interfaces**.
3. Select a data interface on the passive firewall instance and click **Delete**.
4. In the **Delete Network Interface** window, click **Yes, Delete**.
5. Repeat this process for each data interface on the passive firewall instance.

 ©2024 Palo Alto Networks, Inc.

Use Case: Secure the EC2 Instances in the AWS Cloud

In this example, the VPC is deployed in the 10.0.0.0/16 network with two /24 subnets: 10.0.0.0/24 and 10.0.1.0/24. The VM-Series firewall will be launched in the 10.0.0.0/24 subnet to which the internet gateway is attached. The 10.0.1.0/24 subnet is a private subnet that will host the EC2 instances that need to be secured by the VM-Series firewall; any server on this private subnet uses NAT for a routable IP address (which is an Elastic IP address) to access the internet. Use the Planning Worksheet for the VM-Series in the AWS VPC to plan the design within your VPC; recording the subnet ranges, network interfaces and the associated IP addresses for the EC2 instances, and security groups, will make the setup process easier and more efficient.

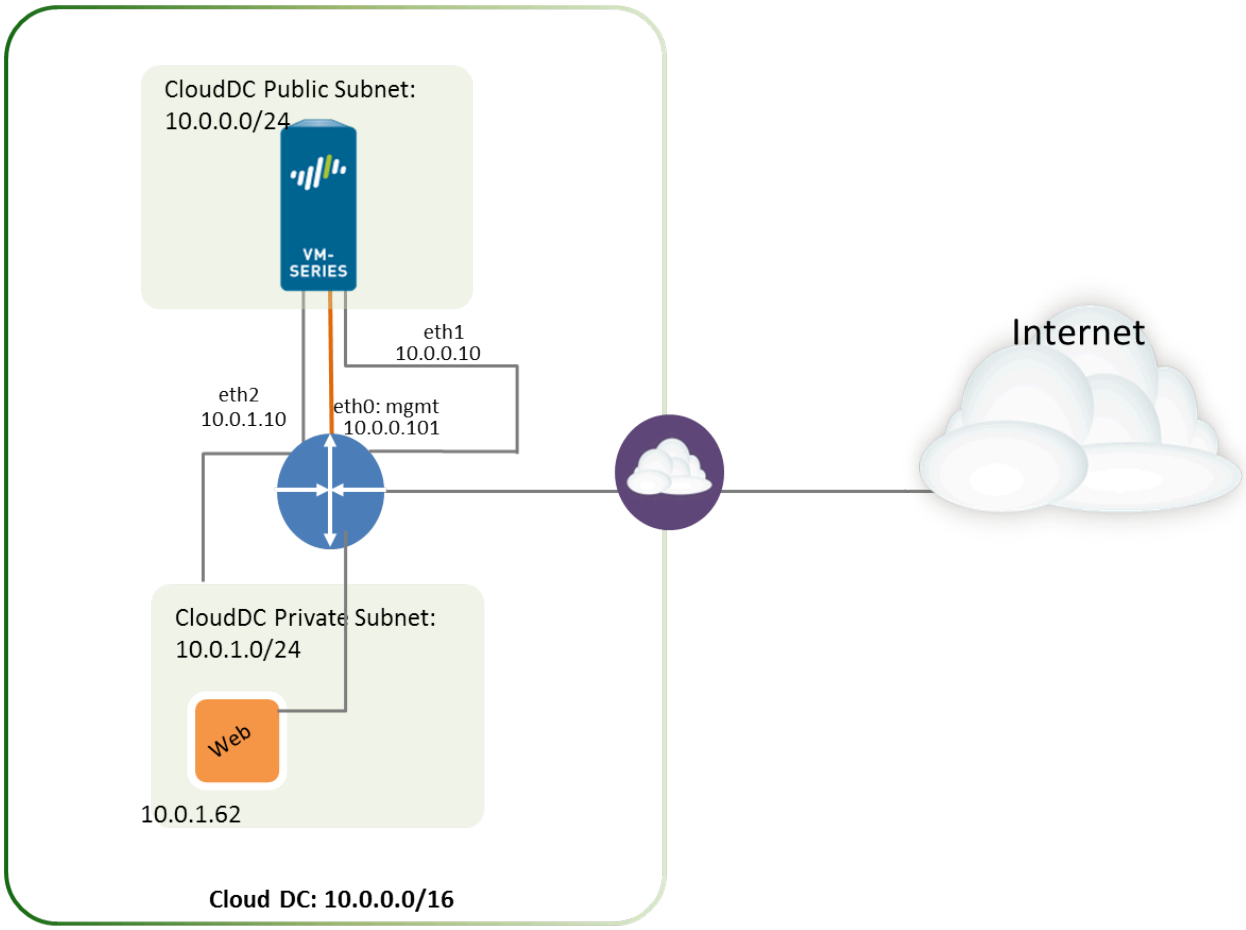

The following image depicts the logical flow of traffic to/from the web server to the internet. Traffic to/from the web server is sent to the data interface of the VM-Series firewall that is attached to the private subnet. The firewall applies policy and processes incoming/outgoing traffic from/to the internet gateway of the VPC. The image also shows the security groups to which the data interfaces are attached.

 ©2024 Palo Alto Networks, Inc.

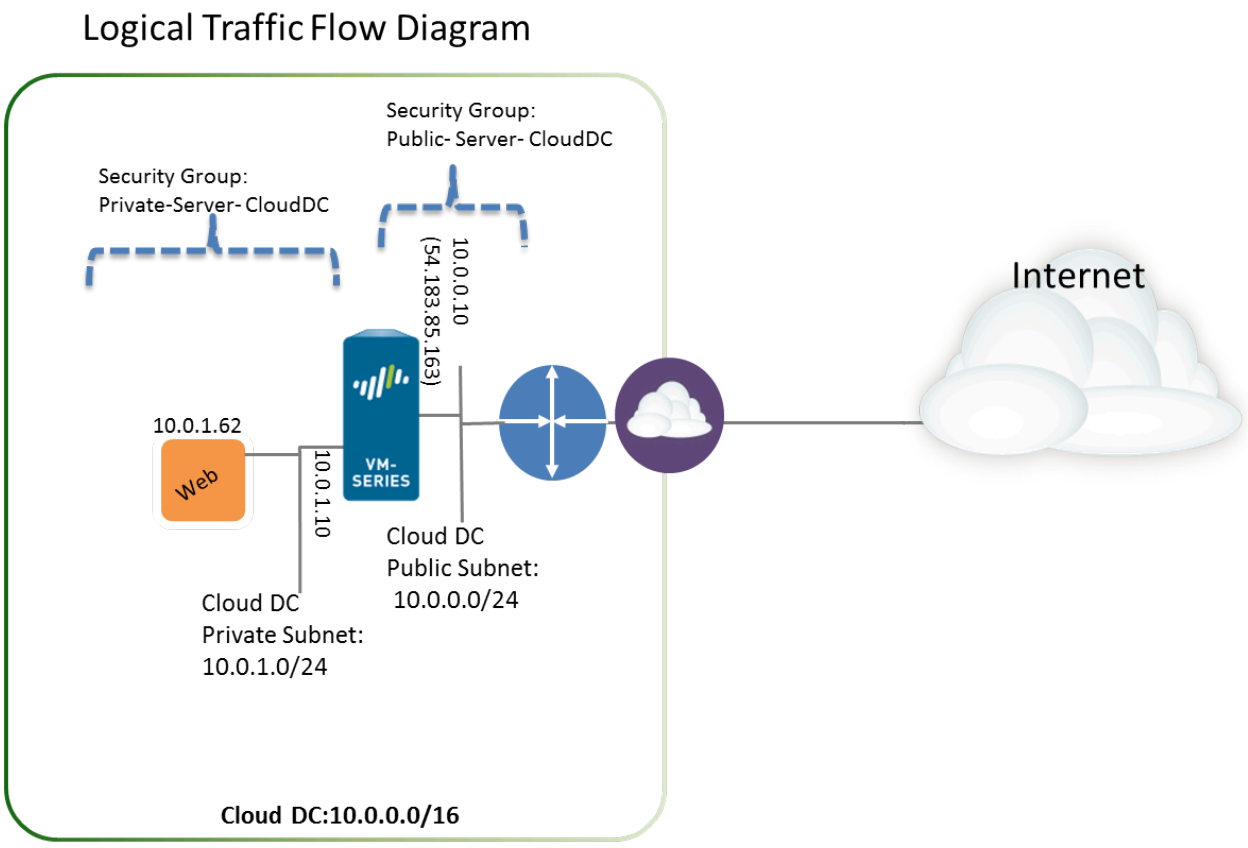

The management IP address and the security group associated with it, are not depicted in this diagram

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create a new VPC with a public subnet (or select an existing VPC).

1. Log in to the AWS console and select the **VPC** Dashboard.
2. Verify that you've selected the correct geographic area (AWS region). The VPC will be deployed in the currently selected region.
3. Select **Start VPC Wizard**, and select **VPC with a Single Public Subnet**.

   In this example, the IP CIDR block for the VPC is 10.0.0.0/16, the VPC name is Cloud DC, the public subnet is 10.0.0.0/24, and the subnet name is Cloud DC Public subnet. You will create a private subnet after creating the VPC.

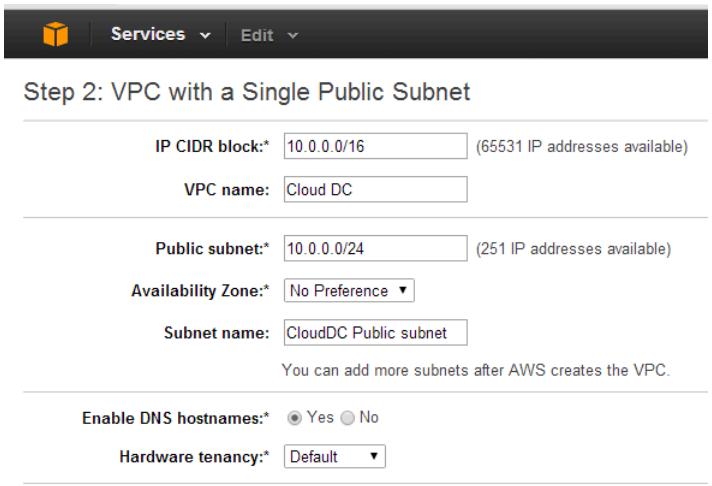

4. Click **Create VPC**.

STEP 2 | Create a private subnet.

Select **Subnets**, and click **Create a Subnet**. Fill in the information.

In this example, the **Name tag** for the subnet is Web/DB Server Subnet, it is created in the Cloud Datacenter VPC and is assigned a CIDR block of 10.0.1.0/24.

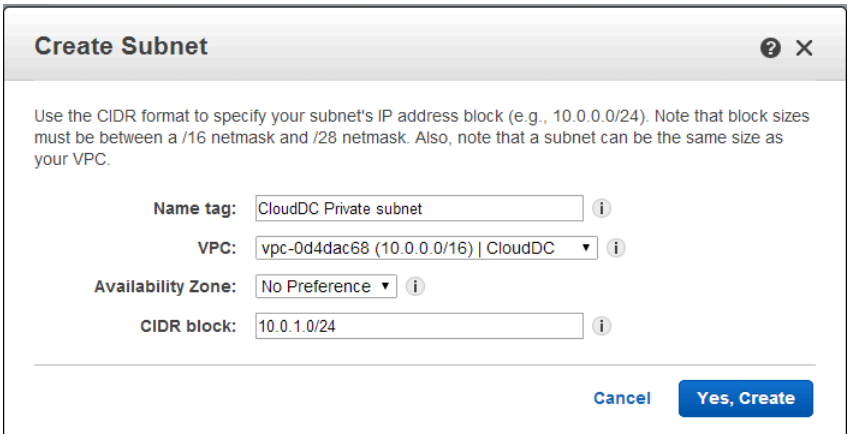

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create a new route table for each subnet.

Although a main route table is automatically created on the VPC, we recommend creating new route tables instead of modifying the default route table.

To direct outbound traffic from each subnet, you will add routes to the route table associated with each subnet, later in this workflow.

1. Select **Route Tables** > **Create Route Table**.
2. Add a **Name**, for example CloudDC-public-subnet-RT, select the **VPC** you created in Step 1, and click **Yes, Create**.
3. Select the route table, click **Subnet Associations** and select the public subnet.

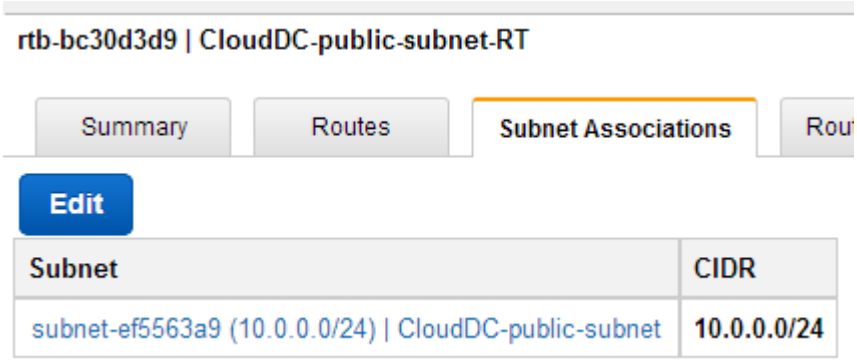

4. Select **Create Route Table**.
5. Add a **Name**, for example CloudDC-private-subnet-RT, select the **VPC** you created in Step 1, and click **Yes, Create**.
6. Select the route table, click **Subnet Associations** and select the private subnet.

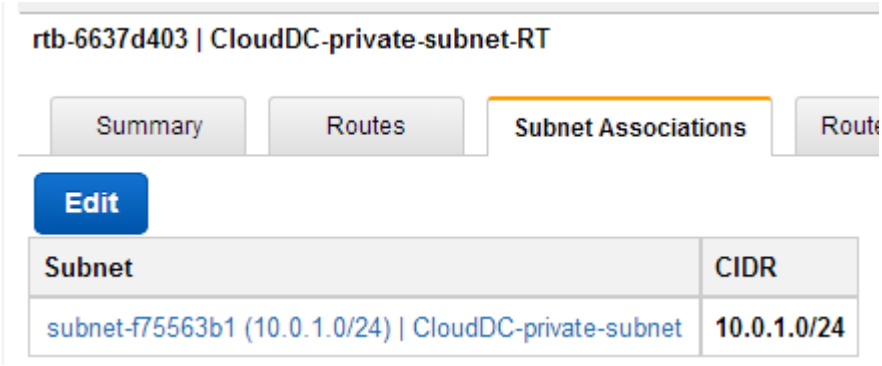

STEP 4 | Create Security Groups to restrict inbound/outbound internet access to the EC2 instances in the VPC.

By default, AWS disallows communication between interfaces that do not belong to the same security group.

Select **Security Groups** and click the **Create Security Group** button. In this example, we create three security groups with the following rules for inbound access:

- CloudDC-Management that specifies the protocols and source IP addresses that can connect to the management interface of the VM-Series firewall. At a minimum you need

 ©2024 Palo Alto Networks, Inc.

SSH, and HTTPS. In this example, we enable SSH, ICMP, HTTP, and HTTPS on the network interfaces that are attached to this security group.

The management interface (eth 0/0) of the VM-Series firewall will be assigned to CloudDC-management-sg.

- Public-Server-CloudDC that specifies the source IP addresses that can connect over HTTP, FTP, SSH within the VPC. This group allows traffic from the external network to the firewall.

  The dataplane interface eth1/1 of the VM-Series firewall will be assigned to Public-Server-CloudDC.

- Private-Server-CloudDC that has very limited access. It only allows other EC2 instances on the same subnet to communicate with each other, and with the VM-Series firewall.

  The dataplane interface eth1/2 of the VM-Series firewall and the application in the private subnet will be attached to this security group.

  The following screenshot shows the security groups for this use case.

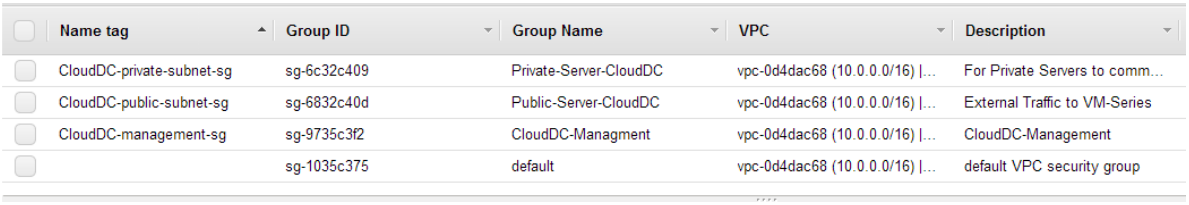

Name tag	Group ID	Group Name	VPC	Description
CloudDC-private-subnet-sg	sg-6c32c409	Private-Server-CloudDC	vpc-0d4dac68 (10.0.0.0/16) |...	For Private Servers to comm...
CloudDC-public-subnet-sg	sg-6832c40d	Public-Server-CloudDC	vpc-0d4dac68 (10.0.0.0/16) |...	External Traffic to VM-Series
CloudDC-management-sg	sg-9735c3f2	CloudDC-Managment	vpc-0d4dac68 (10.0.0.0/16) |...	CloudDC-Management
	sg-1035c375	default	vpc-0d4dac68 (10.0.0.0/16) |...	default VPC security group

STEP 5 | Deploy the VM-Series firewall.

Only the primary network interface that will serve as the management interface will be attached and configured for the firewall during the initial launch. The network interfaces required for handling data traffic will be added in Step 6.

See Step 3 in Launch the VM-Series Firewall on AWS.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Create and attach virtual network interface(s), referred to as Elastic Network Interfaces (ENIs), to the VM-Series firewall. These ENIs are used for handling data traffic to/from the firewall.

1. On the EC2 Dashboard, select **Network Interfaces**, and click **Create Network Interface**.
2. Enter a descriptive name for the interface.
3. Select the subnet. Use the subnet ID to make sure that you have selected the correct subnet. You can only attach an ENI to an instance in the same subnet.
4. Enter the **Private IP** address that you want to assign to the interface or select **Auto-assign** to automatically assign an IP address within the available IP addresses in the selected subnet.
5. Select the **Security group** to control access to the network interface.
6. Click **Yes, Create**.

   In this example, we create two interfaces with the following configuration:

Name	Network interfa	Subnet ID	VPC ID	Zone	Security group	Description	Instance ID
CloudDC-VM-Series-Untrust	eni-bcf355e5	subnet-ef5563a9	vpc-0d4dac68	us-west-1a	Public-Server-...	CloudDC-VM-Series-untrust	i-a7358ff9
CloudDC-VM-Series-Trust	eni-abf355f2	subnet-f75563b1	vpc-0d4dac68	us-west-1a	Private-Server-...	CloudDC-VM-Series-Trust	i-a7358ff9

   - For Eth1/1 (VM-Series-Untrust)
     - Subnet: 10.0.0.0/24
     - Private IP:10.0.0.10
     - Security group: Public-Server-CloudDC
   - For Eth1/2 (VM-Series-Trust)
     - Subnet: 10.0.1.0/24
     - Private IP: 10.0.1.10
     - Security group: Private-Server-CloudDC
7. To attach the ENI to the VM-Series firewall, select the interface you just created, and click **Attach**.

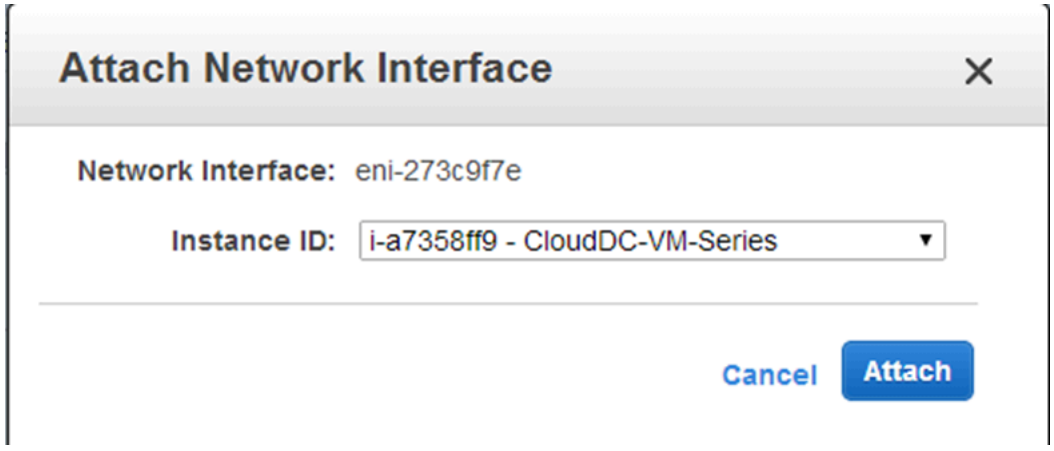

8. Select the **Instance ID** of the VM-Series firewall, and click **Attach**.

 ©2024 Palo Alto Networks, Inc.

9. Repeat steps 7 and 8 to attach the other network interface.

STEP 7 | Create an Elastic IP address and attach it to the firewall dataplane network interface that requires direct internet access.

In this example, VM-Series_Untrust is assigned an EIP. The EIP associated with the interface is the publicly accessible IP address for the web server in the private subnet.

1. Select **Elastic IPs** and click **Allocate New Address**.
2. Select **EC2-VPC** and click **Yes, Allocate**.
3. Select the newly allocated EIP and click **Associate Address**.
4. Select the **Network Interface** and the **Private IP address** associated with the interface and click **Yes, Associate**.

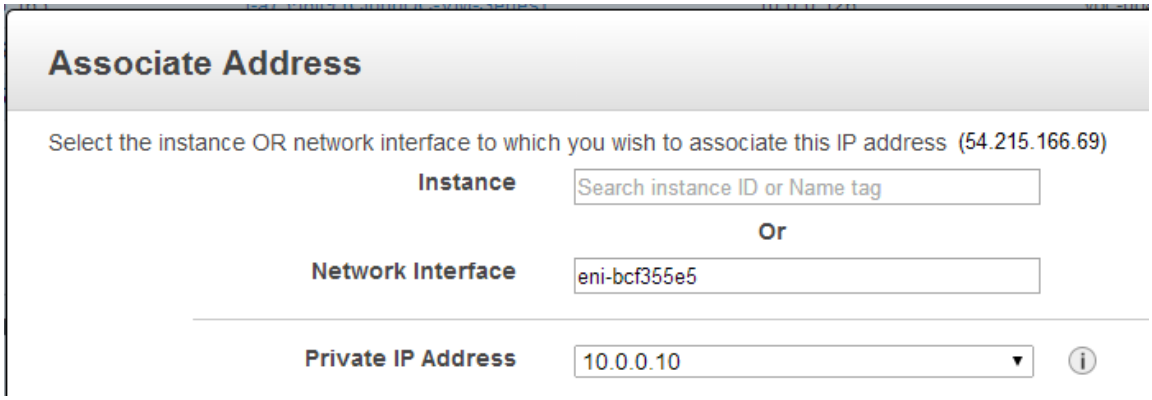

In this example, the configuration is:

Address	Instance	Private IP Address	Scope	Public DNS
54.183.85.163	i-a7358ff9 (CloudDC-VM-Series)	10.0.0.126	vpc-0d4dac68	ec2-54-183-85-163.us-west-...
54.215.166.69	i-a7358ff9 (CloudDC-VM-Series)	10.0.0.10	vpc-0d4dac68	ec2-54-215-166-69.us-west-...

STEP 8 | Disable Source/Destination check on each network interface attached to the VM-Series firewall. Disabling this attribute allows the interface to handle network traffic that is not destined to its IP address.

1. Select the network interface in the **Network Interfaces** tab.
2. In the **Action** drop-down, select **Change Source/Dest. Check**.
3. Click **Disabled** and **Save** your changes.
4. Repeat steps 1-3 for additional network interfaces, firewall-1/2 in this example.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | In the route table associated with the public subnet (from step 3), add a default route to the internet gateway for the VPC.

1. From the VPC Dashboard, select **Route Tables** and find the route table associated with the public subnet.
2. Select the route table, select **Routes** and click **Edit**.
3. Add a route to forward packets from this subnet to the internet gateway. In this example, 0.0.0.0 indicates that all traffic from/to this subnet will use the internet gateway attached to the VPC.

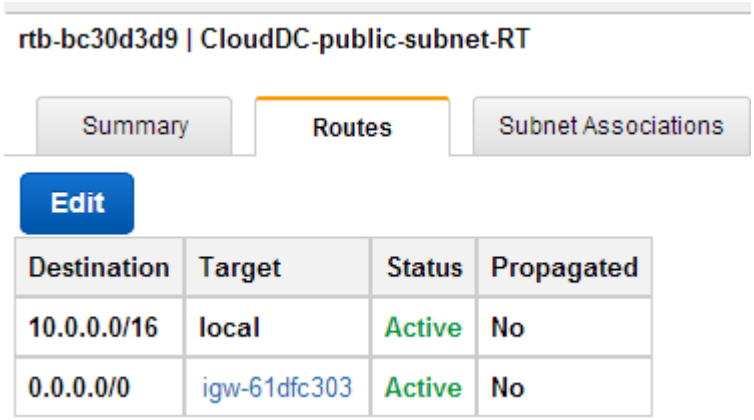

STEP 10 | In the route table associated with the private subnet, add a default route to send traffic to the VM-Series firewall.

Adding this route enables the forwarding of traffic from the EC2 instances in this private subnet to the VM-Series firewall.

1. From the VPC Dashboard, select **Route Tables** and find the route table associated with the private subnet.
2. Select the route table, select **Routes** and click **Edit**.
3. Add a route to forward packets from this subnet to the VM-Series firewall network interface that resides on the same subnet. In this example, 0.0.0.0/0 indicates that all

 ©2024 Palo Alto Networks, Inc.

traffic from/to this subnet will use eni-abf355f2 (ethernet 1/2, which is CloudDC-VM-Series-Trust) on the VM-Series firewall.

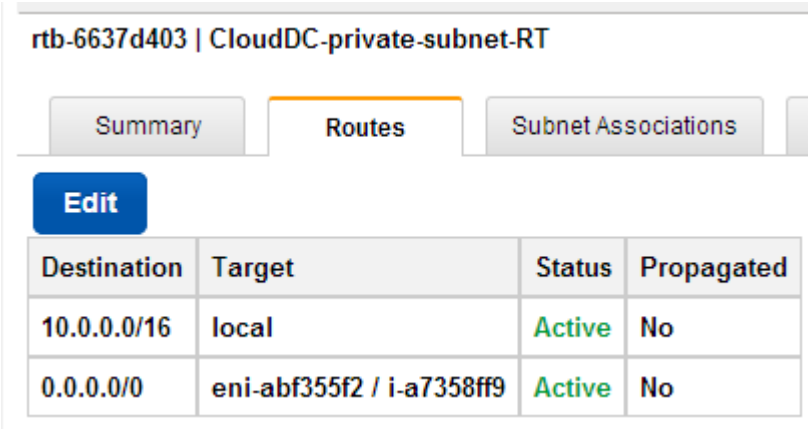

For each web or database server deployed on an EC2 instance in the private subnet, you must define a default route to the IP address of the VM-Series firewall so that the firewall is the default gateway for the server.

Perform steps 11 through 16 on the VM-Series firewall.

STEP 11 | Configure a new administrative password for the firewall.

An SSH tool such as PuTTY is required to access the CLI on the firewall and change the default administrative password. You cannot access the web interface until you SSH and change the default password.

1. Use the public IP address you configured on the firewall, to SSH into the Command Line Interface (CLI) of the VM-Series firewall.

   You will need the private key that you used or created in Launch the VM-Series Firewall on AWS, steps 3-12 to access the CLI.
2. Enter the following command to log in to the firewall:

   ssh -i <*private_key_name*> admin@<*public-ip_address*>
3. Configure a new password, using the following command and follow the onscreen prompts:

   configure

   set mgt-config users admin password

   commit
4. Terminate the SSH session.

STEP 12 | Access the web interface of the VM-Series firewall.

Open a web browser and enter the EIP of the management interface. For example: https://54.183.85.163

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Activate the licenses on the VM-Series firewall. This step is only required for the BYOL license; the usage-based licenses are automatically activated.

See Activate the License.

STEP 14 | On the VM-Series firewall, configure the dataplane network interfaces on the firewall as Layer 3 interfaces.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **Layer3**
   - Select the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select New Zone. Define a new zone, for example untrust, and then click OK.
   - Select IPv4, select **DHCP Client**; the private IP address that you assigned to the network interface in the AWS management console will be acquired automatically.
   - On the **Advanced** > **Other Info** tab, expand the Management Profile drop-down, and select **New Management Profile**.
   - Enter a **Name** for the profile, such as allow_ping, and select **Ping** from the Permitted Services list, then click **OK**.
   - To save the interface configuration, click **OK**.
3. Click the link for **ethernet 1/2** and configure as follows:
   - **Interface Type**: **Layer3**
   - Select the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example trust, and then click **OK**.
   - Select **IPv4**, select **DHCP Client**.
   - On the **IPv4** tab, clear the **Automatically create default route to default gateway provided by server** check box. For an interface that is attached to the private subnet in the VPC, disabling this option ensures that traffic handled by this interface does not flow directly to the IGW on the VPC.
   - On the **Advanced** > **Other Info**, expand the Management Profile drop-down, and select the allow_ping profile you created earlier.
   - Click **OK** to save the interface configuration.
4. Click **Commit** to save the changes. Verify that the Link state for the interface is up . If the link state is not up, reboot the firewall.

Interface	Interface Type	Management Profile	Link State	IP Address	Virtual Router	Security Zone
ethernet1/1	Layer3	allow_ping		Dynamic-DHCP Client	default	untrust
ethernet1/2	Layer3	allow_ping		Dynamic-DHCP Client	default	trust

 ©2024 Palo Alto Networks, Inc.

STEP 15 | On the VM-Series firewall, create Destination NAT and Source NAT rules to allow inbound/outbound traffic to/from the applications deployed within the VPC.

1. Select **Policies** > **NAT**.
2. Create a Destination NAT rule that steers traffic from the firewall to the web server.
   1. Click **Add**, and enter a name for the rule. For example, NAT2WebServer.
   2. In the **Original Packet** tab, make the following selections:
      - **Source Zone**: untrust (where the traffic originates)
      - **Destination Zone**: untrust (the zone for the firewall dataplane interface with which the EIP for the web server is associated.)
      - **Source Address**: Any
      - **Destination Address**: 10.0.0.10
      - In the **Translated Packet** tab, select the Destination Address Translation check box and set the **Translated Address**: to 10.0.1.62, which is the private IP address of the web server.
   3. Click **OK**.

			Original Packet						Translated Packet	
	Name	Tags	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1	NAT2WebServer	none	untrust	untrust	any	any	10.0.0.10	any	none	address: 10.0.1.62

3. Create a Source NAT rule to allow outbound traffic from the web server to the internet.
   1. Click **Add**, and enter a name for the rule. For example, NAT2External.
   2. In the **Original Packet** tab, make the following selections:
      - **Source Zone**: trust (where the traffic originates)
      - **Destination Zone**: untrust (the zone for the firewall dataplane interface with which the EIP for the web server is associated.)
      - **Source Address**: Any
      - **Destination Address**: Any
   3. In the **Translated Packet** tab, make the following selections in the Source Address Translation section:
      - **Translation Type**: Dynamic IP and Port
      - **Address Type**: Translated Address
      - **Translated Address**: 10.0.0.10 (the firewall dataplane interface in the untrust zone.)
   4. Click **OK**.

 ©2024 Palo Alto Networks, Inc.

			Original Packet						Translated Packet	
	Name	Tags	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation
1	NAT2WebServer	none	untrust	untrust	any	any	10.0.0.10	any	none	address: 10.0.1.62
2	NAT2External	none	trust	untrust	any	any	any	any	dynamic-ip-and-port 10.0.0.10	none

4. Click **Commit** to save the NAT policies.

STEP 16 | On the VM-Series firewall, create security policies to manage traffic.

Instead of entering a static IP address for the web server, use a dynamic address group. Dynamic address groups allow you to create policy that automatically adapts to changes so that you do not need to update the policy when you launch additional web servers in the subnet. For details, see Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC.

1. Select **Policies** > **Security**.

   In this example, we have four rules. A rule that allows management access to the firewall traffic, a rule to allow inbound traffic to the web server, a third rule to allow internet

 ©2024 Palo Alto Networks, Inc.

access to the web server, and in the last rule we modify a predefined intrazone-default rule to log all traffic that is denied.

2. Create a rule to allow management access to the firewall.
   1. Click **Add** and enter a **Name** for the rule. Verify that the **Rule Type** is universal.
   2. In the **Source** tab, add untrust as the **Source Zone**.
   3. In the **Destination** tab, add trust as the **Destination Zone**.
   4. In the **Applications** tab, **Add** ping and ssh.
   5. In the **Actions** tab, set the **Action** to Allow.
   6. Click **OK**.

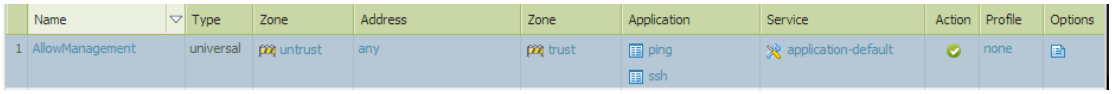

3. Create a rule to allow inbound traffic to the web server.
   1. Click **Add** and enter a **Name** for the rule and verify that the **Rule Type** is universal.
   2. In the **Source** tab, add untrust as the **Source Zone**.
   3. In the **Destination** tab, add trust as the **Destination Zone**.
   4. In the **Applications** tab, **Add** web-browsing.
   5. In the **Service/URL Category** tab, verify that the service is set to application-default.
   6. In the **Actions** tab, set the **Action** to Allow.
   7. In the Profile Settings section of the **Actions** tab, select **Profiles** and then attach the default profiles for antivirus, anti-spyware, and vulnerability protection.
   8. Click **OK**.

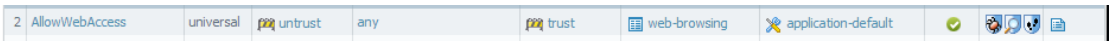

4. Create a rule to allow internet access to the web server.
   1. Click **Add** and enter a **Name** for the rule and verify that the Rule Type is universal.
   2. In the **Source** tab, add trust as the **Source Zone**.
   3. In the Source Address section of the **Source** tab, add 10.0.1.62, the IP address of the web server.
   4. In the **Destination** tab, add untrust as the **Destination Zone**.
   5. In the **Service/URL Category** tab, verify that the service is set to **application-default**.
   6. In the **Actions** tab, set the **Action** to Allow.
   7. In the Profile Settings section of the **Actions** tab, select **Profiles** and then attach the default profiles for antivirus, anti-spyware, and vulnerability protection.
   8. Click **OK**.

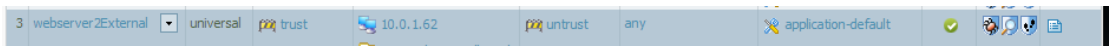

 ©2024 Palo Alto Networks, Inc.

5. Edit the interzone-default rule to log all traffic that is denied. This predefined interzone rule is evaluated when no other rule is explicitly defined to match traffic across different zones.
   1. Select the **interzone-default** rule and click **Override**.
   2. In the **Actions** tab, select **Log at session end**.
   3. Click **OK**.

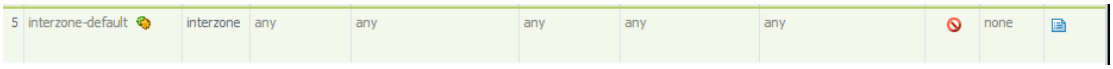

6. Review the complete set of security rules defined on the firewall.
7. Click **Commit** to save the policies.

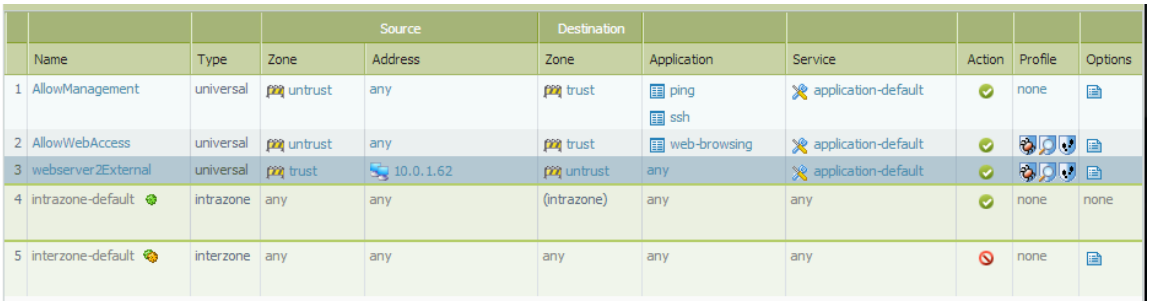

			Source		Destination					
	Name	Type	Zone	Address	Zone	Application	Service	Action	Profile	Options
1	AllowManagement	universal	untrust	any	trust	ping ssh	application-default		none	
2	AllowWebAccess	universal	untrust	any	trust	web-browsing	application-default			
3	webserver2External	universal	trust	10.0.1.62	untrust	any	application-default			
4	intrazone-default	intrazone	any	any	(intrazone)	any	any		none	none
5	interzone-default	interzone	any	any	any	any	any		none	

 ©2024 Palo Alto Networks, Inc.

STEP 17 | Verify that the VM-Series firewall is securing traffic.

1. Launch a web browser and enter the IP address for the web server.
2. Log in to the web interface of the VM-Series firewall and verify that you can see the traffic logs for the sessions at **Monitor** > **Logs** > **Traffic**.

   - Traffic inbound to the web server (arrives at EC2 instance in the AWS VPC):

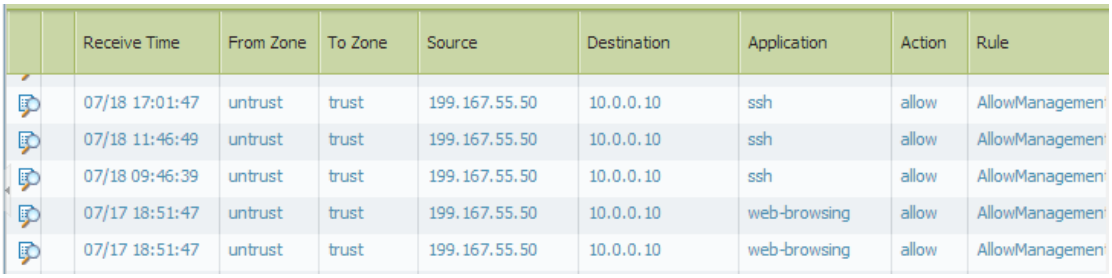

Receive Time	From Zone	To Zone	Source	Destination	Application	Action	Rule
07/18 17:01:47	untrust	trust	199.167.55.50	10.0.0.10	ssh	allow	AllowManagemen
07/18 11:46:49	untrust	trust	199.167.55.50	10.0.0.10	ssh	allow	AllowManagemen
07/18 09:46:39	untrust	trust	199.167.55.50	10.0.0.10	ssh	allow	AllowManagemen
07/17 18:51:47	untrust	trust	199.167.55.50	10.0.0.10	web-browsing	allow	AllowManagemen
07/17 18:51:47	untrust	trust	199.167.55.50	10.0.0.10	web-browsing	allow	AllowManagemen

   - Traffic outbound from the web server (EC2 instance in the AWS VPC):

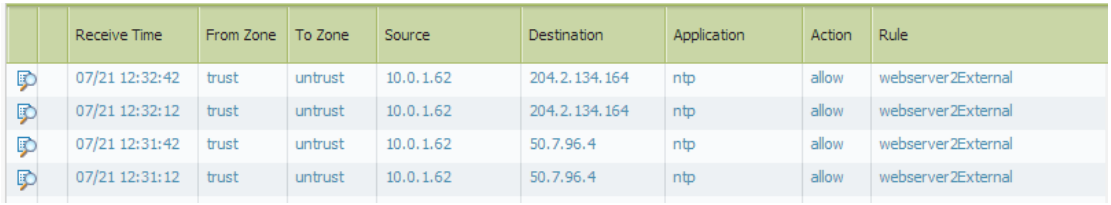

Receive Time	From Zone	To Zone	Source	Destination	Application	Action	Rule
07/21 12:32:42	trust	untrust	10.0.1.62	204.2.134.164	ntp	allow	webserver2External
07/21 12:32:12	trust	untrust	10.0.1.62	204.2.134.164	ntp	allow	webserver2External
07/21 12:31:42	trust	untrust	10.0.1.62	50.7.96.4	ntp	allow	webserver2External
07/21 12:31:12	trust	untrust	10.0.1.62	50.7.96.4	ntp	allow	webserver2External

You have successfully deployed the VM-Series firewall as a cloud gateway!

 ©2024 Palo Alto Networks, Inc.

Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC

In a dynamic environment such as the AWS-VPC where you launch new EC2 instances on demand, the administrative overhead in managing security policy can be cumbersome. Using Dynamic Address Groups in security policy allows for agility and prevents disruption in services or gaps in protection.

In this example, you can use the VM Information Source on the firewall to monitor a VPC and use Dynamic Address Groups in security policy to discover and secure EC2 instances. As you spin up EC2 instances, the Dynamic Address Group collates the IP addresses of all instances that match the criteria defined for group membership, and then security policy is applied for the group. The security policy in this example allows internet access to all members of the group.

Instead of using VM Information Source on the firewall, you can opt to use Panorama as the central point for communicating with your VPCs. Using the AWS plugin on Panorama, you can retrieve the IP address-to-tag mapping and register the information on the managed firewalls for which you configure notification. For more details on this option, see VM Monitoring with the AWS Plugin on Panorama.

This workflow in the following section assumes that you have created the AWS VPC and deployed the VM-Series firewall and some applications on EC2 instances. For instructions on setting up the VPC for the VM-Series, see Use Case: Secure the EC2 Instances in the AWS Cloud.

STEP 1 | Configure the firewall to monitor the VPC.

1. Select **Device** > **VM Information Sources**.
2. Click **Add** and enter the following information:
   1. A **Name** to identify the VPC that you want to monitor. For example, VPC-CloudDC.
   2. Set the **Type** to AWS VPC.
   3. In **Source**, enter the URI for the VPC. The syntax is **ec2.<*your_region*>.amazonaws.com**
   4. Add the credentials required for the firewall to digitally sign API calls made to the AWS services. You need the following:
      - **Access Key ID**: Enter the alphanumeric text string that uniquely identifies the user who owns or is authorized to access the AWS account.
      - **Secret Access Key**: Enter the password and confirm your entry.
   5. (**Optional**) Modify the **Update interval** to a value between 5-600 seconds. By default, the firewall polls every 5 seconds. The API calls are queued and retrieved within

 ©2024 Palo Alto Networks, Inc.

every 60 seconds, so updates may take up to 60 seconds plus the configured polling interval.

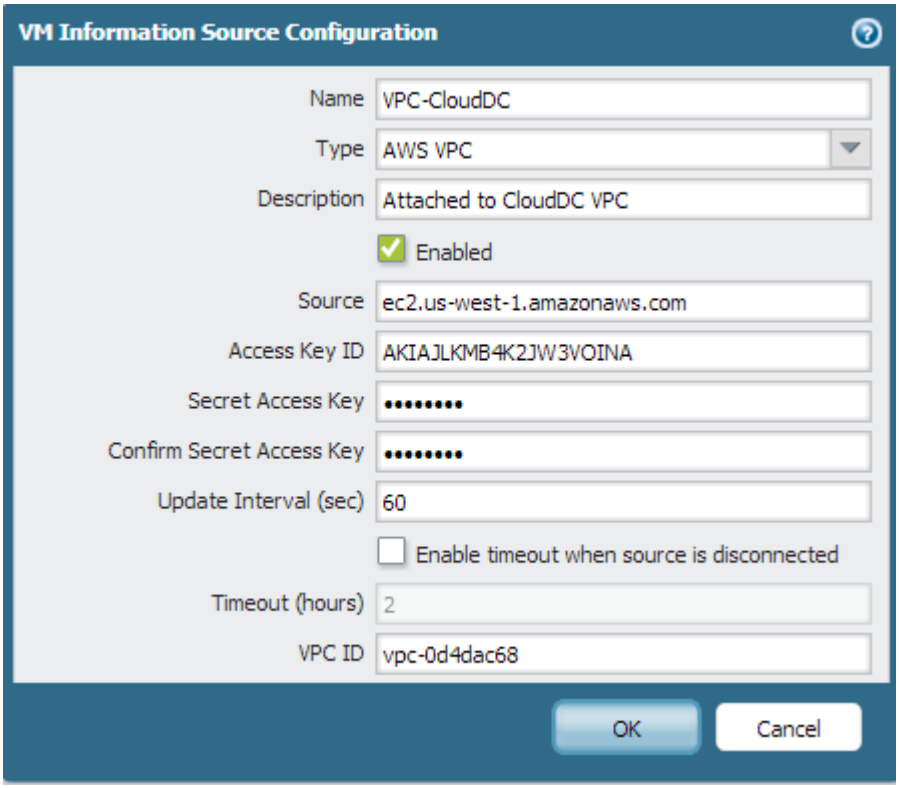

6. Enter the **VPC ID** that is displayed on the VPC Dashboard in the AWS management console.
7. Click **OK**, and **Commit** the changes.
8. Verify that the connection **Status** displays as connected

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Tag the EC2 instances in the VPC.

For a list of tags that the VM-Series firewall can monitor, see List of Attributes Monitored on the AWS VPC.

A tag is a name-value pair. You can tag the EC2 instances either on the EC2 Dashboard on the AWS management console or using the AWS API or AWS CLI.

In this example, we use the EC2 Dashboard to add the tag:

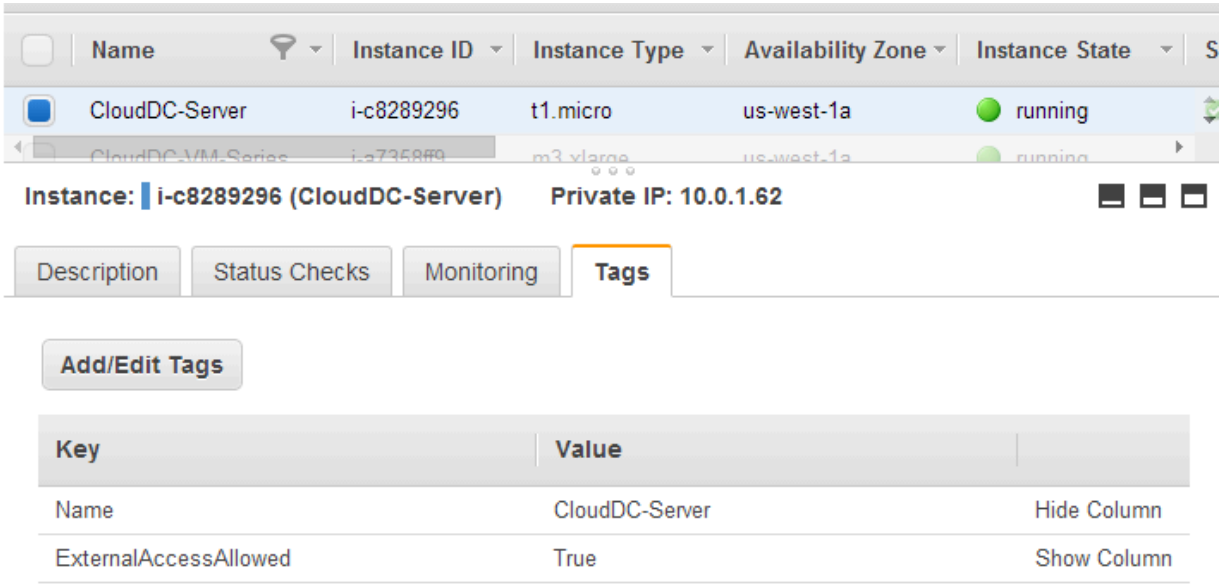

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create a dynamic address group on the firewall.

View the tutorial to see a big picture view of the feature.

1. Select **Object** > **Address Groups**.
2. Click **Add** and enter a **Name** and a **Description** for the address group.
3. Select **Type** as **Dynamic**.
4. Define the match criteria.
   1. Click **Add Match Criteria**, and select the **And** operator.
   2. Select the attributes to filter for or match against. In this example, we select the ExternalAccessAllowed tag that you just created and the subnet ID for the private subnet of the VPC.

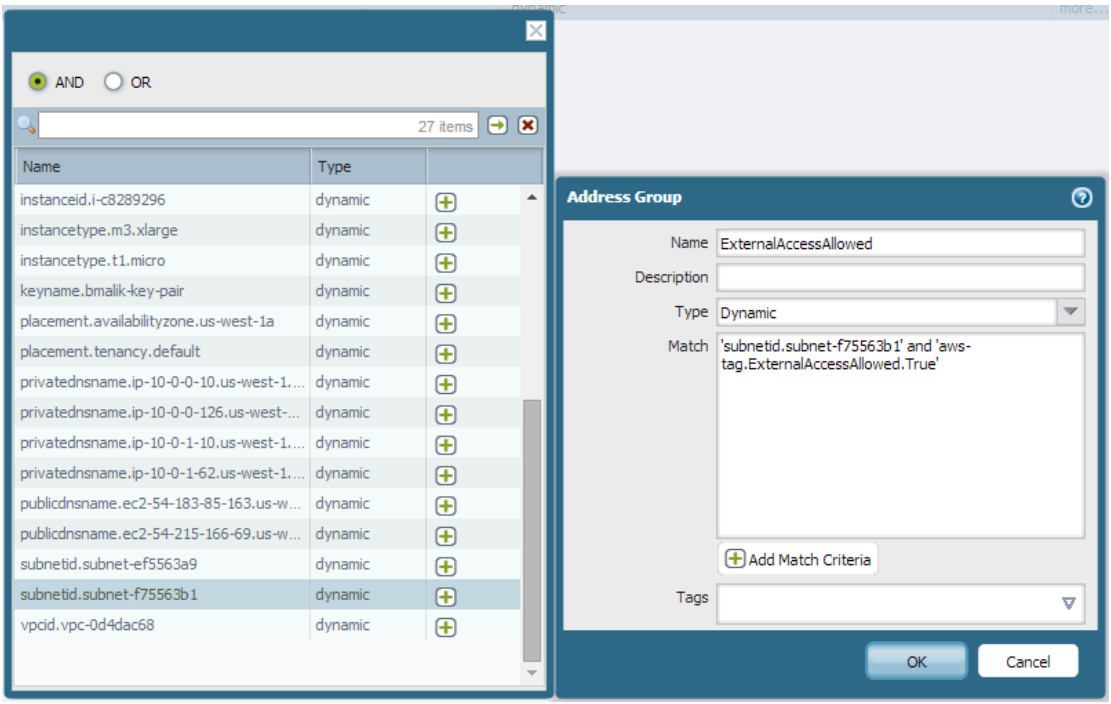

5. Click **OK**.
6. Click **Commit**.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Use the dynamic address group in a security policy.

To create a rule to allow internet access to any web server that belongs to the dynamic address group called ExternalServerAccess.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** for the rule and verify that the **Rule Type** is universal.
3. In the **Source** tab, add trust as the **Source Zone**.
4. In the Source Address section of the **Source** tab, **Add** the ExternalServerAccess group you just created.
5. In the **Destination** tab, add untrust as the **Destination Zone**.
6. In the **Service/URL Category** tab, verify that the service is set to **application-default**.
7. In the **Actions** tab, set the **Action** to Allow.
8. In the Profile Settings section of the **Actions** tab, select **Profiles** and then attach the default profiles for antivirus, anti-spyware, and vulnerability protection.
9. Click **OK**.

10. Click **Commit**.

STEP 5 | Verify that members of the dynamic address group are populated on the firewall.

Policy will be enforced for all IP addresses that belong to this address group, and are displayed here.

1. Select **Policies** > **Security**, and select the rule.
2. Select the drop-down arrow next to the address group link, and select **Inspect**. You can also verify that the match criteria is accurate.
3. Click the **more** link and verify that the list of registered IP addresses is displayed.

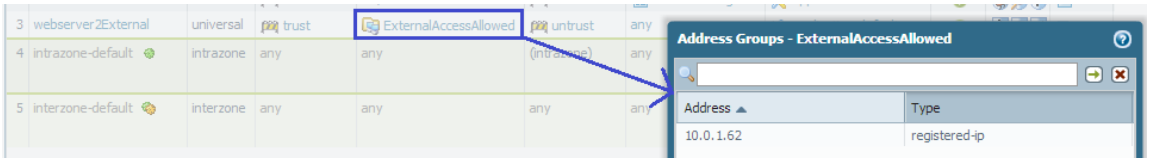

 ©2024 Palo Alto Networks, Inc.

Use Case: VM-Series Firewalls as GlobalProtect Gateways on AWS

Securing mobile users from threats and risky applications is often a complex mix of procuring and setting up the security and IT infrastructure, ensuring bandwidth and uptime requirements in multiple locations around the globe while staying within your budget.

The VM-Series firewall on AWS melds the security and IT logistics required to consistently and reliably protect devices used by mobile users in regions where you do not have a presence. By deploying the VM-Series firewall in the AWS cloud, you can quickly and easily deploy GlobalProtect™ gateways in any region without the expense or IT logistics that are typically required to set up this infrastructure using your own resources.

To minimize latency, select AWS regions that are closest to your users, deploy the VM-Series firewalls on EC2 instances, and configure the firewalls as GlobalProtect gateways. With this solution, the GlobalProtect gateways in the AWS cloud enforce security policy for internet traffic so there is no need to backhaul that traffic to the corporate network. Additionally, for access to resources on the corporate network, the VM-Series firewalls on AWS leverage the LSVPN functionality to establish IPSec tunnels back to the firewall on the corporate network.

For ease of deployment and centralized management of this distributed infrastructure, use Panorama to configure the GlobalProtect components used in this solution. Optionally, to ensure that mobile devices, such as smartphones and tablets, are safe for use on your network, use a Mobile Device Manager to configure and manage mobile devices.

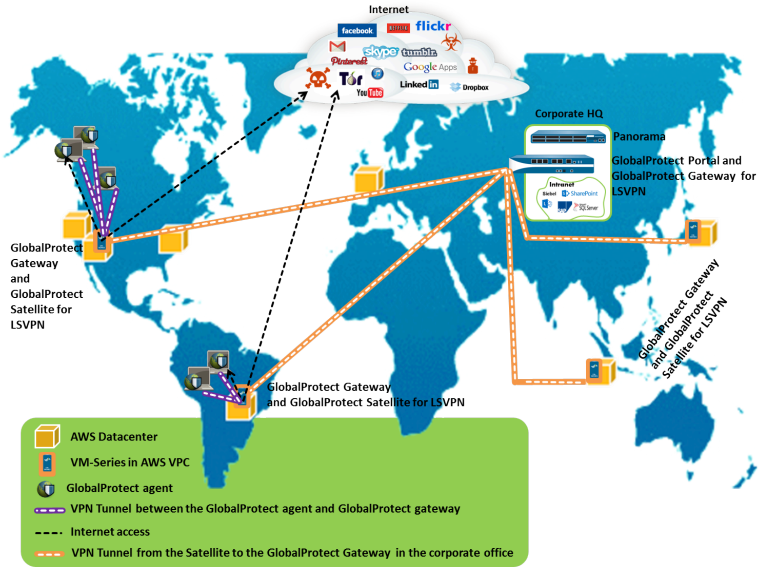

- Components of the GlobalProtect Infrastructure
- Deploy GlobalProtect Gateways on AWS

Components of the GlobalProtect Infrastructure

To block risky applications and protect mobile users from malware, you must set up the GlobalProtect infrastructure, which includes the GlobalProtect portal, the GlobalProtect gateway,

 ©2024 Palo Alto Networks, Inc.

and the GlobalProtect app. Additionally, for access to corporate resources, you must set up an IPSec VPN connection between the VM-Series firewalls on AWS and the firewall in the corporate headquarters using LSVPN (a hub and spoke VPN deployment).

- The GlobalProtect agent/app is installed on each end-user system that is allowed to access corporate applications and resources. The agent first connects to the portal to obtain information on the gateways and then establishes a secure VPN connection to the closest GlobalProtect gateway. The VPN connection between the end-user system and the gateway ensures data privacy.
- The GlobalProtect portal provides the management functions for the GlobalProtect infrastructure. Every end-user system receives configuration information from the portal, including information about available gateways as well as any client certificates that may be required to connect to the GlobalProtect gateway(s). In this use case, the GlobalProtect portal is a hardware-based firewall that is deployed in the corporate headquarters.
- The GlobalProtect gateway delivers mobile threat prevention and policy enforcement based on applications, users, content, device, and device state. In this use case, the VM-Series firewalls on AWS function as the GlobalProtect gateways. The GlobalProtect gateway scans each user request for malware and other threats, and, if policy allows, sends the request to the internet or to the corporate network over the IPSec tunnel (to the LSVPN gateway).
- For LSVPN, you must configure the GlobalProtect portal, GlobalProtect gateway for LSVPN (hub), and the GlobalProtect Satellites (spokes).

  In this use case, the hardware-based firewall in the corporate office is deployed as the GlobalProtect portal and the LSVPN gateway. The VM-Series firewalls on AWS are configured to function as GlobalProtect satellites. The GlobalProtect satellites and gateway are configured to establish an IPSec tunnel that terminates on the gateway. When a mobile user requests an application or resource that resides on the corporate network, the VM-Series firewall routes the request over the IPSec tunnel.

Deploy GlobalProtect Gateways on AWS

To secure mobile users, in addition to deploying and configuring the GlobalProtect gateways on AWS, you need to set up the other components required for this integrated solution. The following table includes the recommended workflow:

- Deploy the VM-Series firewall(s) on AWS.

  See Deploy the VM-Series Firewall on AWS.

- Configure the firewall at the corporate headquarters.

  In this use case, the firewall is configured as the GlobalProtect portal and the LSVPN gateway.

  - Configure the GlobalProtectportal.
  - Configure the GlobalProtectportal for LSVPN.
  - Configure the portal to authenticateLSVPN satellites.
  - Configure the GlobalProtectgateway for LSVPN.

 ©2024 Palo Alto Networks, Inc.

- Set up a template on Panorama for configuring the VM-Series firewalls on AWS as GlobalProtect gateways and LSVPN satellites.

  To easily manage this distributed deployment, use Panorama to configure the firewalls on AWS.

  - Create template(s) on Panorama.

  Then use the following links to define the configuration in the templates.

  - Configure the firewall asa GlobalProtect gateway.
  - Prepare the satellite tojoin the LSVPN.

- Create device groups on Panorama to define the network access policies and internet access rules and apply them to the firewalls on AWS.

  See Create device groups.

- Apply the templates and the device groups to the VM-Series firewalls on AWS, and verify that the firewalls are configured properly.

- Deploy the GlobalProtect client software.

  Every end-user system requires the GlobalProtect agent or app to connect to the GlobalProtect gateway.

  See Deploy the GlobalProtectclient software.

 ©2024 Palo Alto Networks, Inc.

Resource Monitoring on AWS

As you deploy or terminate resources in the AWS public cloud, you can either use the Panorama plugin for AWS or use the AWS resource information sources on the firewall to consistently enforce security policy rules on these workloads. See the Compatibility Matrix for Panorama plugin version information.

The Panorama plugin for AWS is built for scale and allows you to monitor up to 1000 AWS VPCs on the AWS public cloud. With this plugin, you use Panorama as an anchor to poll your AWS accounts for tags, and then distribute the metadata (IP address-to-tag mapping) to many firewalls in a device group. Because Panorama communicates with your AWS accounts to retrieve AWS resource information, you're able to streamline the number of API calls made to the cloud environment. When using Panorama and the AWS plugin, you can centralize the retrieval of tags and Security policy management to ensure consistent policies for hybrid and cloud-native architectures. See AWS Resource Monitoring with the AWS Plugin on Panorama.

If you do not have Panorama or you have a simpler deployment and need to monitor 10 VPCs or fewer, you can use the VM Information Source on the firewall (hardware or VM-Series firewall) to monitor your AWS workloads. You can use the metadata, which the firewall retrieves, in Dynamic Address Groups and reference them in Security policies to secure your VM workloads as they spin up or down and IP addresses change frequently. See Use Case: Use Dynamic Address Groups to Secure New EC2 Instances within the VPC.

AWS Resource Monitoring with the AWS Plugin on Panorama

As you deploy or terminate resources in the AWS public cloud, you need a way to synchronously update Security policy on your Palo Alto Networks® firewall(s) so that you can secure these EC2 instances. To enable this capability from Panorama, you must install the AWS plugin on Panorama and enable API communication between Panorama and your AWS VPCs. Panorama can then collect a predefined set of attributes (or metadata elements) as tags for your AWS resources and register the information to your Palo Alto Networks® firewall(s). When you reference these tags in Dynamic Address Groups and match against them in Security policy rules, you can consistently enforce policy across all assets deployed within your AWS accounts.

- Set Up the AWS Plugin for Monitoring on Panorama
- List of Attributes Monitored on the AWS VPC

Set Up the AWS Plugin for VM Monitoring on Panorama

To find all the virtual machine workloads that your organization has deployed in the AWS public cloud, you need to install the AWS plugin on Panorama and configure *Monitoring Definitions* that enable Panorama to authenticate to your AWS VPC(s) and retrieve VM information on the workloads. Panorama retrieves the IP address of the VMs that are running— public IP address, and primary and secondary private IP addresses—and the associated tags. For a list of the metadata elements that Panorama supports, see List of Attributes Monitored on the AWS VPC.

After Panorama fetches the attributes, to push the virtual machine information from Panorama to the firewalls, you must add the firewalls (hardware or VM-Series) as managed devices on Panorama, and group the firewalls into one or more Device Groups. You can then specify which

 ©2024 Palo Alto Networks, Inc.

device groups are part of the *Notify Group*, which is a configuration element in a Monitoring Definition, that Panorama uses to register the IP address-to-tag mapping it retrieves from AWS.

Finally, to consistently enforce Security policies across the EC2 instances, you must set up Dynamic Address Groups and reference them in policy rules that allow or deny traffic to the IP addresses of the VMs. For streamlining your configuration and managing policies and objects centrally from Panorama, you can define the Dynamic Address Groups and Security policy rules on Panorama and push them to the firewalls instead of managing the Dynamic Address Groups and Security policy rules locally on each firewall.

The AWS plugin version 3.0.1 or later is for monitoring EC2 instances for up to 1000 VPCs on the AWS public cloud, AWS GovCloud, and AWS China. However, because Panorama cannot be deployed on AWS China, the IAM role does not support instance profiles on AWS China; you must provide the AWS credentials.

- Planning Checklist for VM Monitoring on AWS
- IAM Roles and Permissions for Panorama
- Install or Upgrade the AWS Plugin
- Configure the AWS Plugin for VM Monitoring

Planning Checklist for VM Monitoring on AWS

For Panorama to interact with the AWS APIs and collect information on your EC2 instances, you need to create an IAM role and assign the policies that grant the permissions required to authenticate to AWS and access the EC2 instances within your VPC. You can add 100 IAM roles to manage up to 1000 VPCs on Panorama.

- ☐ Gather the VPC ID.
- ☐ Tag your EC2 instances on AWS. You can tag (define a name-value pair) the EC2 instances either on the EC2 Dashboard on the AWS management console or using the AWS API or AWS CLI. See List of Attributes Monitored on the AWS VPC for the list of supported attributes.
- ☐ Check for duplicate IP addresses across the VPCs for which you will enable monitoring. If you have duplicate IP addresses across AWS VPCs, the metadata will be appended together or swapped and this may cause unexpected results in policy enforcement.

  Duplicate IP addresses are written to the plugin_aws_ret.log file that you can access from the CLI on Panorama.
- ☐ Review the requirements for Panorama and the managed firewalls:
  - Minimum system requirements—Panorama virtual appliance or hardware-based Panorama appliance.

Panorama Minimum Requirements				
System Resources	Memory	CPUs	Number of Monitored VPCs	Number of Tags Registered
	16GB	4	1-100	Panorama 9.1 or later with AWS plugin v 2.0 is tested to retrieve 10,000 IP addresses with 13 tags

 ©2024 Palo Alto Networks, Inc.

Panorama Minimum Requirements				
	32 GB	8	100-500	for each, or 5000 IP addresses with 25 tags for each, and successfully register them to the firewalls included within a device group.The tag length—includes name and value—for each EC2 instance is assumed to be 64 bytes per tag. For example, the EC2 instance name tag is aws.ec2.tag.Name.prod-web-app-4523-lvss6j.
	64 GB	16	500-1000	
Panorama OS version	10.0.5 or later			
AWS plugin version	3.0.1 or later			
Licenses	Active support license and a device management license on Panorama for managing the firewalls. Next-generation firewalls must also have a valid support license.			
Roles and Permissions to retrieve metadata on the EC2 instances	See IAM Roles and Permissions for Panorama			

- You must add the firewall as managed devices on Panorama and create Device Groups so that you can configure Panorama to notify these groups with the VM information it retrieves. Device groups can include VM-Series firewalls or virtual systems on the hardware firewalls.
- If your Panorama appliances are in a high availability configuration, you must manually install the same version of the AWS plugin on both Panorama peers. Additionally, if you are using instance profiles, you must attach the same instance profile to both Panorama peers.

  You configure the AWS plugin on the active Panorama peer only. On commit, the configuration is synced to the passive Panorama peer. Only the active Panorama peer polls the AWS accounts you have configured for VM Monitoring.
- Set up the credentials/permissions that Panorama requires to digitally sign API calls to the AWS services.

  You can choose whether you want to provide the long-term credentials—Access Key ID and Secret Access Key—that enable access to the resources within each AWS account, or set up an Assume Role on AWS to allow access to defined AWS resources within the same AWS account or cross-accounts. With an Assume Role, you must set up a trust relationship and define the permissions while creating the role itself. This is specifically useful in a cross-

 ©2024 Palo Alto Networks, Inc.

account deployment where the querying account does not have permissions to see or handle data from the queried account. For the Panorama plugin to successfully authenticate to the VPC and retrieve the tags, you must configure the Assume Role to use the AWS Security Token Service (STS) API to any AWS service. And a user from the querying account must have STS permissions to query the Assume Role and obtain the temporary security credentials to access resources. If your Panorama is deployed on AWS, you can opt to use an instance profile instead of providing the AWS credentials for the IAM role. The instance profile includes the role information and associated credentials that Panorama needs to digitally sign API calls to the AWS services. See IAM Roles and Permissions for Panorama for more details.

IAM Roles and Permissions for Panorama

With the AWS plugin version 2.0, you can use IAM roles or instance profiles to enable Panorama to authenticate and retrieve metadata on the resources deployed within your AWS account(s).

- When your Panorama is not deployed on AWS, you have two options. You can either provide the long-term IAM credentials for the AWS accounts you want to monitor, or set up an Assume Role on AWS to allow access to defined AWS resources within the same AWS account or cross-accounts. An Assume Role is recommended as the more secure option.
- When your Panorama is deployed on AWS, in addition to the two options listed above, you can also add an instance profile that allows the IAM role to be passed to the EC2 instance. You can use an instance profile where all your monitored resources and Panorama are hosted within the same account, or an instance profile with Assume Role for cross account access where your Panorama and monitored resources are deployed across different AWS accounts. If you use the instance profile, you do not enter your AWS credentials on Panorama.

Option 1: IAM role with long term credentials	
Roles and Permissions Required	The AWS credentials associated with the AWS account that has the VPC/EC2 instances you want to monitor. The JSON format for the minimum permissions associated with the IAM role with long-term credentials is as follows:

```
{ “Version”: “2012-10-17",
  “Statement”: [
        {
            “Sid”: “VisualEditor0”,
            “Effect”: “Allow”,
            “Action”: [

 “elasticloadbalancing:DescribeLoadBalancerAttributes”,

 “elasticloadbalancing:DescribeLoadBalancers”,
                “elasticloadbalancing:DescribeTags”,
                “ec2:DescribeInstances”,
                “ec2:DescribeNetworkInterfaces”,
                “ec2:DescribeVpcs”,
                “ec2:DescribeVpcEndpoints”,
                “ec2:DescribeSubnets”
            ],
            “Resource”: “*”
```

 ©2024 Palo Alto Networks, Inc.

```
        }
    ]
}
```

Inputs on Panorama	Enter the **Access Key ID** and **Secret Access Key** for the user in **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Role**.

Option 2: IAM role with Assume Role

Roles and Permissions Required

While you can use this option to monitor VPCs within the same or cross account, this option is recommended to enable cross account access by assuming a role that allows you to access resources to which you may normally have access.

To assume a role from a different account, your AWS account must be trusted by that role and defined as a trusted entity in its trust policy. In addition, a user who wants to access a role in a different account must have a policy with secure token service (STS) access that specifies the role ARN.

On Account 1 that you want to monitor:

- Create an IAM role with required permissions. For VM Monitoring you need the following permissions.

```
{ “Version”: “2012-10-17",
  “Statement”: [
        {
            “Sid”: “VisualEditor0”,
            “Effect”: “Allow”,
            “Action”: [

 “elasticloadbalancing:DescribeLoadBalancerAttributes”,

 “elasticloadbalancing:DescribeLoadBalancers”,
                “elasticloadbalancing:DescribeTags”,
                “ec2:DescribeInstances”,
                “ec2:DescribeNetworkInterfaces”,
                “ec2:DescribeVpcs”,
                “ec2:DescribeVpcEndpoints”,
                “ec2:DescribeSubnets”
            ],
            “Resource”: “*”
        }
    ]
}
```

- Copy the Role ARN.
- Create a user and add the Account ID for Account 2 as a trusted entity. This allows Account 2 the permissions to use this role to access the resources within your Account 1.

On Account 2 that requires access to account 1

 ©2024 Palo Alto Networks, Inc.

	• Attach the following policy with STS permissions and modify the Role ARN to match what you created on Account 1. `{ "Version": "2012-10-17", "Statement": { "Effect": "Allow", "Action": "sts:AssumeRole", "Resource":"arn:aws:iam::012347211234:role/ PAN-OS-assume-role" } }`
Inputs on Panorama	• Enter the **Access Key ID** and **Secret Access Key** for the user on Account 2 on **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Role**. • Enter the **Role ARN** for the AWS Account 1 which you want to monitor in the **Panorama** > **Plugins** > **AWS** > **Monitoring Definitions**.
Option 3: Instance profile	
Roles and Permissions Required	Only when Panorama is deployed as an EC2 instance on AWS *Note that when you use the AWS Management console to create an IAM role, the console automatically creates an instance profile with the same name as the role. Because the role and the instance profile has the same name, when you launch tour Panorama (EC2 instance) with an IAM role, the instance profile of the same name is associated with it.* When Panorama and the resources you want to monitor are all in a single AWS account. Create an IAM role with AmazonEC2ReadOnlyAccess.
Inputs on Panorama	Select **Instance Profile** as the option in **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Role**.
Option 4: Instance profile with Assume Role	
Roles and Permissions Required	Use instance profile with Assume role when Panorama and the resources you want to monitor are deployed across AWS accounts. For Panorama HA, make sure to attach the same instance profile to both Panorama peers. **On Account 1, where your EC2 instances are deployed:** • Create an IAM role. • To this role, add the AWS Account ID (Account 2) where your Panorama is deployed as a trusted entity. • Attach the JSON policies as detailed above for VM Monitoring.

 ©2024 Palo Alto Networks, Inc.

	• Copy the Role ARN.This role is required for Panorama to retrieve metadata on your EC2 instances or EKS clusters. **On Account 2, where your Panorama is deployed:** • Create an IAM role and attach the JSON policy (with the STS policy and resource ARN you got from Account 1). • For each additional AWS account you want to monitor, copy the same STS policy and modify the Role ARN.
Inputs on Panorama	• Select **Instance Profile** as the option in **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Role** • Enter the **Role ARN** for the AWS account which you want to monitor in the **Panorama** > **Plugins** > **AWS** > **Monitoring Definitions**. For example Account 1 in this example.

Install or Upgrade the AWS Plugin

To get started with monitoring your EC2 instances on AWS, refer to the Compatibility Matrix for the Panorama Plugin for AWS and VM-Series plugin versionsrequired to support VM monitoring.

To upgrade a Panorama plugin for AWS version 3.0.1, you must first upgrade the plugin version to 2.0.2.

After you install the AWS plugin version 3.0.1, you cannot downgrade to version 2.0.x or below.

If you have a Panorama HA configuration, repeat the installation/upgrade process on each Panorama peer.

Install or uninstall plugins during a planned maintenance window.

If you currently have a Panorama plugin for any cloud platform installed, installing (or uninstalling) an additional plugin requires a Panorama reboot so commit changes.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away

 ©2024 Palo Alto Networks, Inc.

- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Log in to the Panorama Web Interface, select **Panorama** > **Plugins** and click **Check Now** to get the **AWS** plugin version that supports VM monitoring.

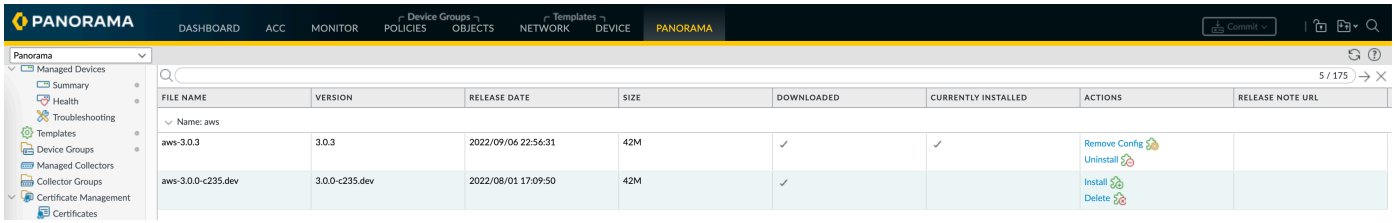

STEP 2 | **Download and Install** the plugin.

After you successfully install, Panorama refreshes and the AWS plugin displays on the **Panorama** > **Plugins** tab.

*On the Panorama **Dashboard** General Information widget you can verify the Panorama Plugin for AWS version that is installed.*

STEP 3 | (Panorama in HA) **Commit** > **Commit to Panorama**.

If your Panorama is in HA, commit the changes to the Panorama configuration to ensure that tags are registered to the Panorama peer on failover.

Configure the AWS Plugin for VM Monitoring

To begin monitoring the virtual machines in your AWS public cloud deployment, after you Install the AWS Plugin you must create a Monitoring Definition. This definition specifies the IAM Role that is authorized to access the EC2 instances within the AWS VPC you want to monitor and the Notify Group that includes the firewalls to which Panorama should push all the IP-address-to-tag mappings it retrieves. In order to enforce policy, you must then create Dynamic Address Groups and reference them in Security policy. The Dynamic Address Groups enable you to filter the tags you want to match on, so that the firewall can get the public and private IP addresses registered against each tag, and then allow or deny access to traffic to and from the workloads based on the policy rules you define.

STEP 1 | Log in to the Panorama web interface.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Set up the following objects for enabling VM Monitoring on AWS.

- Verify that monitoring is enabled on the plugin. This setting must be enabled for Panorama to communicate with the AWS public cloud for VM Monitoring.

  The checkbox for **Enable Monitoring** is on **Panorama** > **Plugins** > **AWS** > **Setup** > **General**.

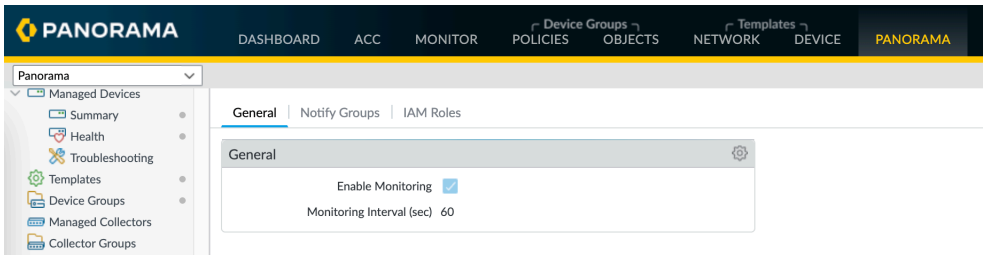

- Add a notify group.

  1. Select **Panorama** > **Plugins** > **AWS** > **Setup** > **Notify Groups** > **Add**.

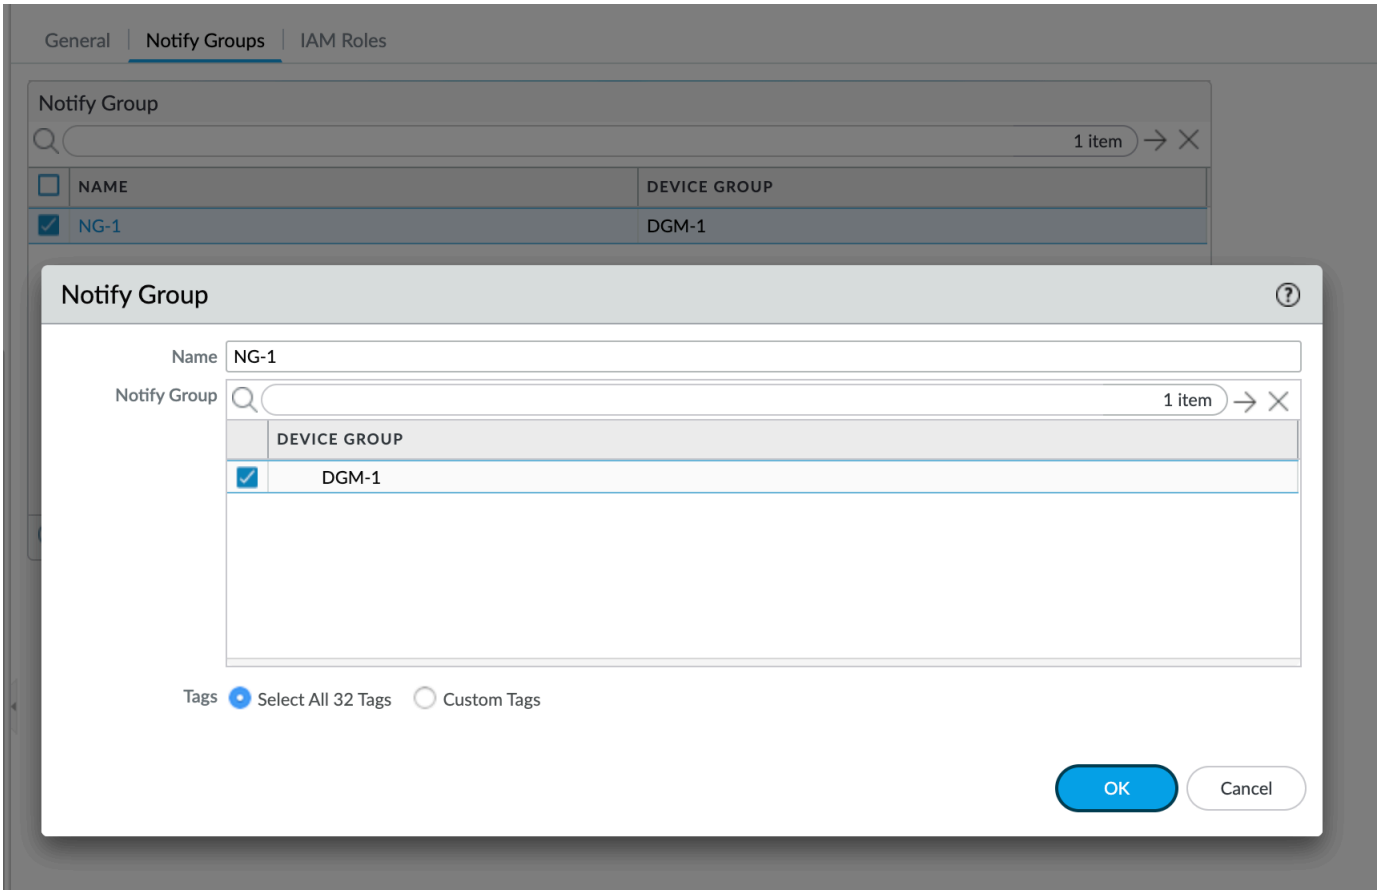

  2. Enter a **Name** to identify the group of firewalls to which Panorama pushes the VM information it retrieves.
  3. Select the **Device Groups**, which are a group of firewalls or virtual systems, to which Panorama will push the VM information (IP address-to-tag mapping) it retrieves from your AWS VPCs. The firewalls use the update to determine the most current list of members that constitute dynamic address groups referenced in policy. If you are using

 ©2024 Palo Alto Networks, Inc.

the Panorama plugin for Azure and AWS, you can target the same firewall or virtual system with tags from both environments.

Think through your Device Groups carefully.

- *Because a Monitoring Definition can include only one notify group, make sure to select all the relevant Device Groups within your notify group. If you want to unregister the tags that Panorama has pushed to a firewall included in a notify group, you must delete the Monitoring Definition.*
- *To register tags to all virtual systems on a firewall enabled for multiple virtual systems, you must add each virtual system to a separate device group on Panorama and assign the device groups to the notify group. If you assign all the virtual systems to one device group, Panorama will register tags to only one virtual system on the firewall.*

4. Select the tags that you want to retrieve from the AWS VPCs.

   You can **Select All 32 Tags** (the default) or pick the **Custom Tags** you want to retrieve for your instances. With the Custom Tags option, you can **Add** the predefined tags and the user-defined tags that you want to use as match criteria in Security policy. If you are monitoring a large number of EC2 instances, reducing the number of tags you retrieve

 ©2024 Palo Alto Networks, Inc.

ensures more efficient use of the CPU and memory capacity on your Panorama. Refer to Planning Checklist for VM Monitoring on AWS for some guidelines.

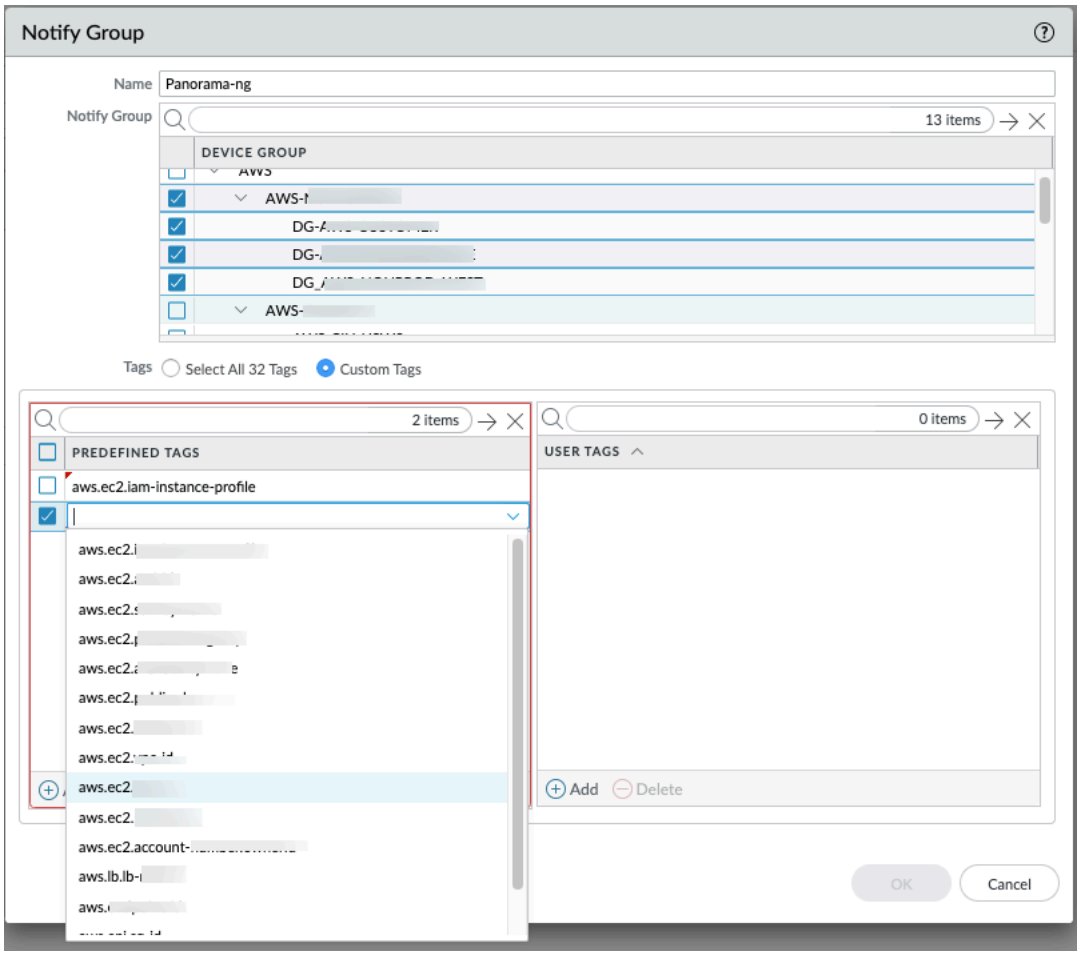

- Add an IAM Role.

  An IAM role is an entity that allows you to delegate access so that Panorama can make service requests on your behalf to the AWS resources (virtual machines that are deployed as EC2 instances).

  1. Select **Panorama** > **Plugins** > **AWS** > **Setup** > **IAM Role** > **Add**.

 ©2024 Palo Alto Networks, Inc.

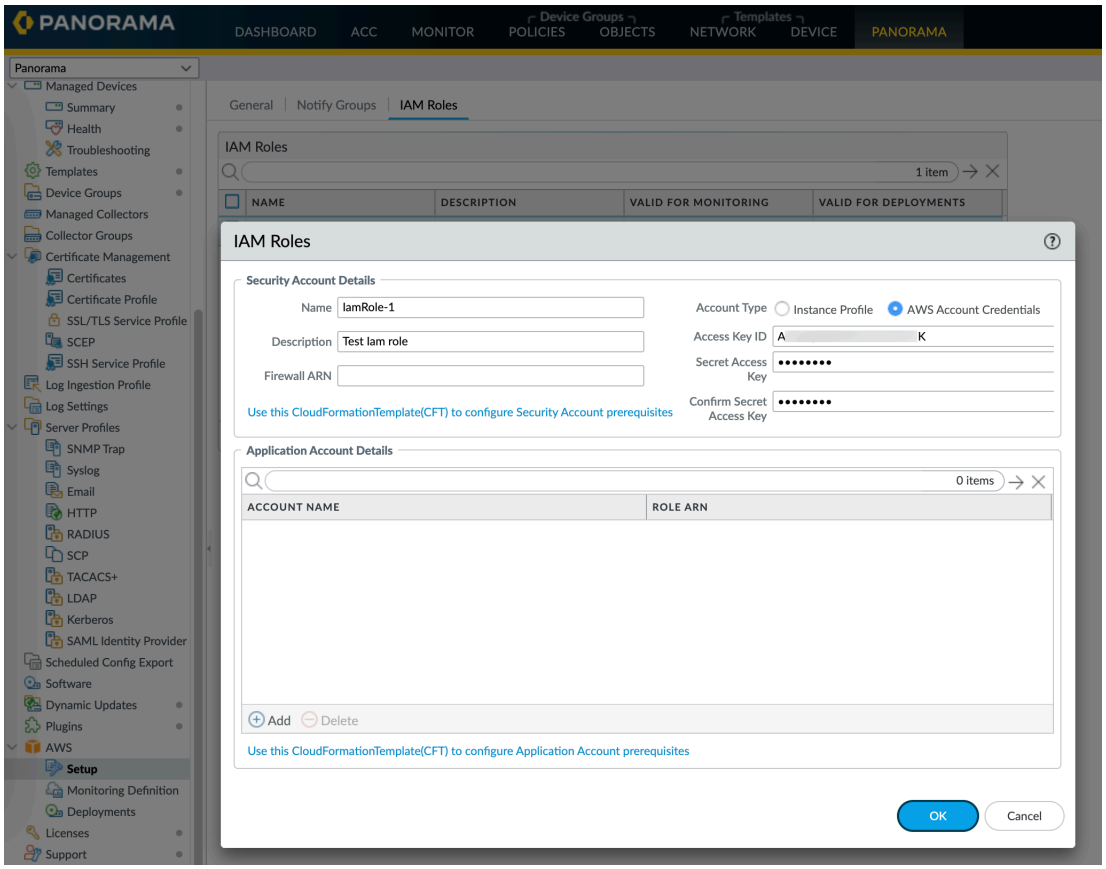

2. Enter a **Name** and optionally a **Description** to identify the IAM role.
3. Select Account Type—**Instance Profile** or **AWS Account Credentials**.If your Panorama is deployed on AWS, you can choose to either attach an instance profile with the correct permissions to your Panorama or add the credentials associated with the IAM role on Panorama. If your Panorama is not deployed on AWS, you must enter the credentials for the IAM role locally on Panorama.
4. (For AWS Account Credentials only) Enter the **Secret Access Key** and re-enter it to confirm, and click **OK**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create a **Monitoring Definition** for each VPC you want to monitor.

When you add a new Monitoring definition, it is enabled by default.

- Select **Panorama** > **Plugins** > **AWS** > **Monitoring Definition** > **General**, to **Add** a new definition.
- Enter a **Name** and optionally a **Description** to identify the AWS VPC for which you use this definition.
- Select the **IAM Role**, **Add** the **VPC ID** from the VPC Dashboard on the AWS management console, and **Notify Group**.
- Select **AWS Regions**:
  - **All**—Select all AWS regions.
  - **Select**—Select specific AWS regions. Search AWS regions from the **Member** search bar or **Add** new regions.
- (Optional) Enter the **Role ARN**, if you have set up role chaining and IAM roles with temporary credentials that have permissions to use the AWS STS API to access AWS

 ©2024 Palo Alto Networks, Inc.

resources with the same account or cross-account. The Role ARN must belong to the VPC you want to monitor.

- Select a **Notify Group**, and **Enable** monitoring.

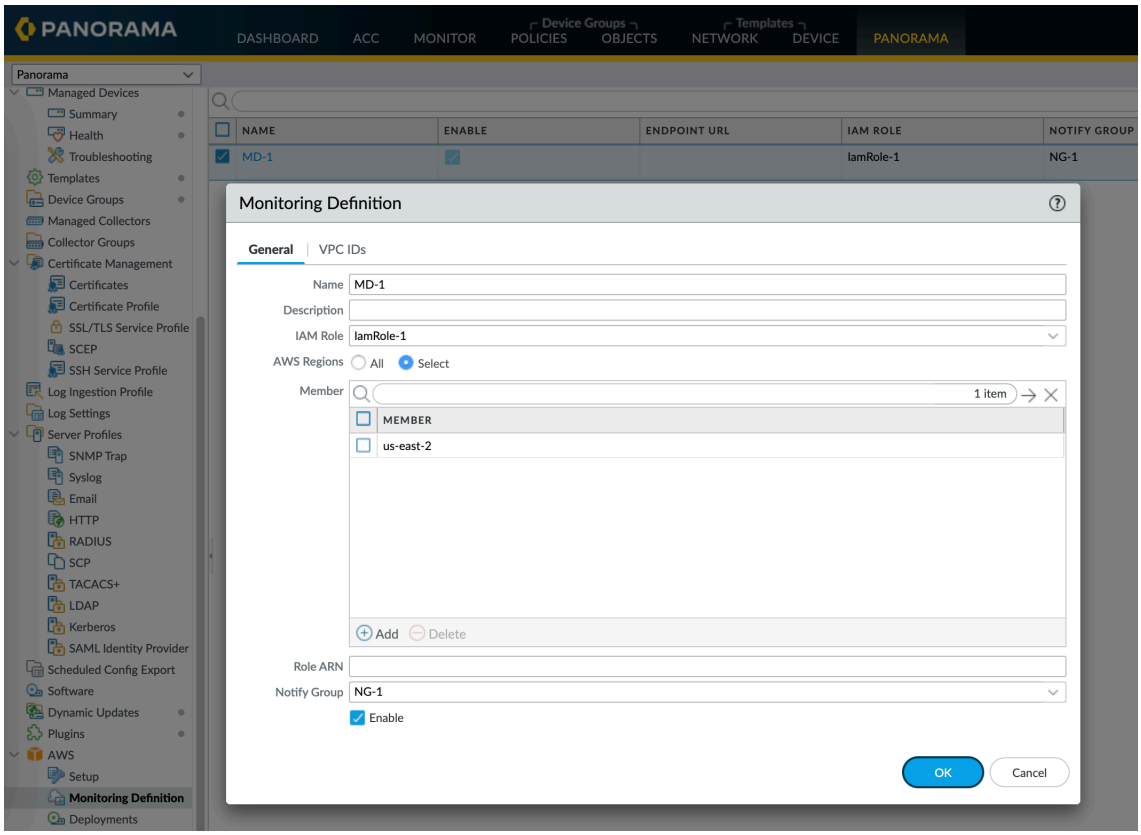

- On the **VPC IDs** tab, add the VPC IDs from the VPC dashboard on the AWS Management Console.

 ©2024 Palo Alto Networks, Inc.

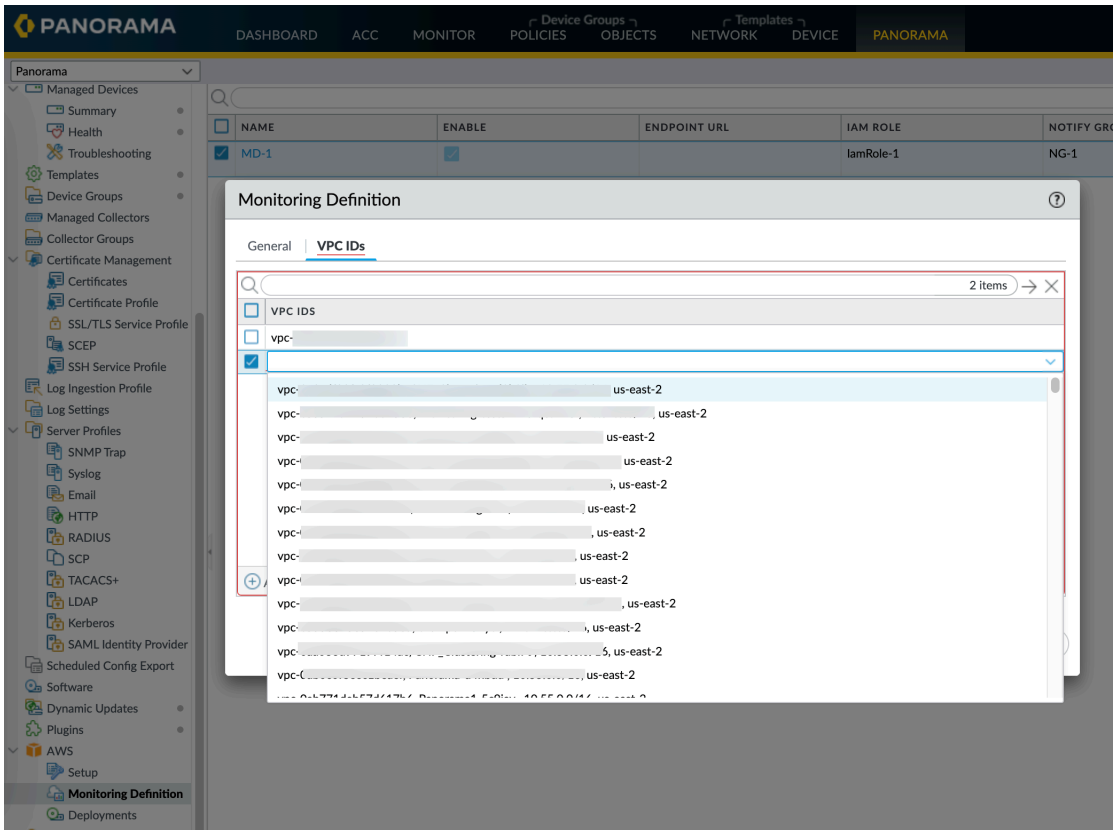

STEP 4 | **Commit** the changes on Panorama.

Verify that the status for the Monitoring Definition displays as Success. If it fails, verify that you entered the AWS VPC ID accurately and provided the correct keys and IDs for authorizing access.

*Click **Validate** to verify that Panorama can authenticate using the IAM role and keys and to communicate with the AWS VPCs you've entered above.*

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Verify that you can view the VM information on Panorama, and define the match criteria for Dynamic Address Groups.

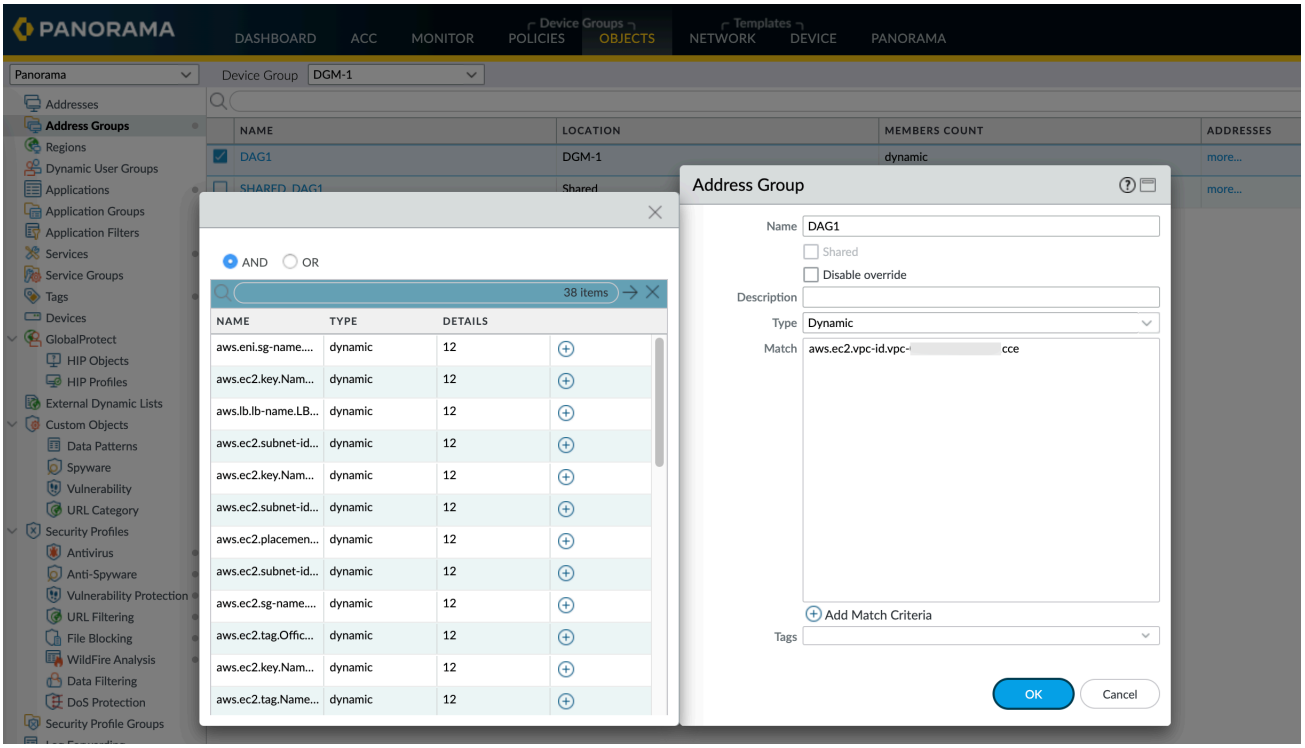

On HA failover, the newly active Panorama attempts to reconnect to the AWS cloud and retrieve tags for all monitoring definitions. If Panorama is unable to reconnect with even one of the monitoring definitions that you have configured and enabled, Panorama generates a system log message

```
Unable to process accounts after HA switch-over; user-intervention required.
```

If this happens, you must log into Panorama and verify the monitoring definitions to fix invalid credentials or remove invalid accounts. Although Panorama is disconnected from the AWS cloud, all tags that were retrieved for the monitoring definitions before the failover, are retained and the firewalls can continue to enforce policy on that list of IP addresses. Panorama removes all tags associated with the accounts only when you delete a monitoring definition. As a best practice, to monitor this issue, you can configure action-oriented log forwarding to an HTTPS destination from Panorama so that you can take action immediately.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Know where to find the logs related to the AWS plugin on Panorama for troubleshooting.

- Use the CLI command **less plugins-log** to view a list of all available logs

  less plugins-log plugin_aws_ret.log displays logs related to IP address and tag retrieval.

  less plugins-log plugin_aws_proc.log displays logs related to processing of the registered IP address and tags.

  less plugins-log plugin_aws.log displays logs related to the AWS plugin configuration and daemons.

Use **show plugins aws vm-mon-status** for the status of the Monitoring Definitions.

```
admin@Panorama> show plugins aws vm-mon-status
Mon-Def Name      VPC           Status  Last Updated Time
 Error Msg
--------------------------------------------------------------------------------
MD-Ins-Prof-ARN  vpc-07986b091  Success 2019-12-02T10:24:56.007000
MD-gov           vpc-7ea1cf1a   Success 2019-12-02T10:24:56.008000
MD-IAM-ARN       vpc-025a83c123 Success 2019-12-02T10:24:56.012000
```

 ©2024 Palo Alto Networks, Inc.

Auto Scaling VM-Series Firewalls with the Amazon ELB Service

The Palo Alto Networks auto scaling templates for AWS help you deploy configure and deploy VM-Series firewalls to protect applications deployed in AWS. The templates leverage AWS scalability features to independently and automatically scale VM-Series firewalls deployed in AWS to meet surges in application workload resource demand.

- VM-Series automation capabilities include the PAN-OS API and bootstrapping (using a bootstrap file for version 2.0, and Panorama for version 2.1).
- AWS automation technology includes CloudFormation templates and scripts for AWS services such as Lambda, auto scaling groups (ASGs), Elastic Load Balancing (ELB), S3, and SNS.

The templates are available on the Palo Alto Networks GitHub repository for Auto Scaling VM-Series Firewalls in AWS:

- Version 2.0 provides a firewall template and an application template. These templates and the supporting scripts deploy VM-Series firewalls, an internet facing firewall, an internal firewall, and application ASGs in a single Virtual Private Cloud (VPC) or multiple VPCs.

  In version 2.0, Palo Alto Networks supports the firewall template while the application template is community-supported. See VM-Series Auto Scale Template for AWS Version 2.0 for deployment details.
- Version 2.1 adds support for deployment in a single VPC and adds support for a load balancer sandwich topology that enables you to deploy the VM-Series firewalls in to a front-end VPC and the back-end applications in to one or more application VPCs connected by VPC peering or AWS PrivateLink.

  In version 2.1 you can implement both application load balancers (ALBs) and network load balancers (NLBs) in VPCs. Version 2.1 includes two firewall templates and five application templates. See VM-Series Auto Scaling Templates for AWS Version 2.1 for deployment details.

If you have an existing template deployment, there is no migration procedure.

The following table compares some high-level features of each template version.

Features / Requirements	Version 2.0	Version 2.1
Panorama running PAN-OS 9.0.1 or a later release in Panorama mode.	(Optional) If you choose to use Panorama, you must configure VPC peering between the VM-Series firewall VPC and the application VPCs.	(Required) Deploy the Version 2.1 templates.

 ©2024 Palo Alto Networks, Inc.

Features / Requirements	Version 2.0	Version 2.1
Panorama in a high availability (HA) configuration is not supported.	Peered traffic traverses the public internet.	*On Panorama, you must manually install the VM-Series plugin to enable VM-Series firewalls to publish PAN-OS metrics for auto scaling.*
Bootstrapping	`bootstrap.xml` config file in an S3 bucket.	An `init-cfg.txt` file for Panorama.
Palo Alto Networks S3 bucket sample	Use your own S3 bucket or use the sample in panw-aws-autoscale-v20-us-west-2.	Use your own S3 bucket for the deployment.
Single VPC or separate VPCs (hub and spoke)	Yes	Yes
New VPC	Yes	Yes
Existing VPC (brown field)	No	Yes
Availability zones per VPC	2	2-4
External load balancer	ALB only	ALB or NLB
Internal load balancer	NLB only	ALB or NLB
AWS PrivateLink connection to the VM-Series firewall VPC and the backend servers.	No	Yes

For details on the templates see:

- VM-Series Auto Scaling Templates for AWS Version 2.0
- VM-Series Auto Scaling Templates for AWS Version 2.1

 ©2024 Palo Alto Networks, Inc.

VM-Series Auto Scaling Templates for AWS Version 2.0

To help you manage increased application scaling, version 2.0 of the auto scaling VM-Series firewall template provides a hub and spoke architecture that simplifies deployment. This version of the solution provides two templates that support a single and multiple-VPC deployment both within a single AWS account and across multiple AWS accounts.

- **Firewall Template**—The firewall template deploys an application load balancer (ALB) and VM-Series firewalls within auto scaling groups across two Availability Zones (AZs). This internet-facing ALB distributes traffic that enters the VPC across a pool of VM-Series firewalls. The VM-Series firewalls automatically publish custom PAN-OS metrics that enable auto scaling.

  Palo Alto Networks officially supports the firewall template and, with a valid support entitlement, you can request assistance from Palo Alto Networks Technical Support.

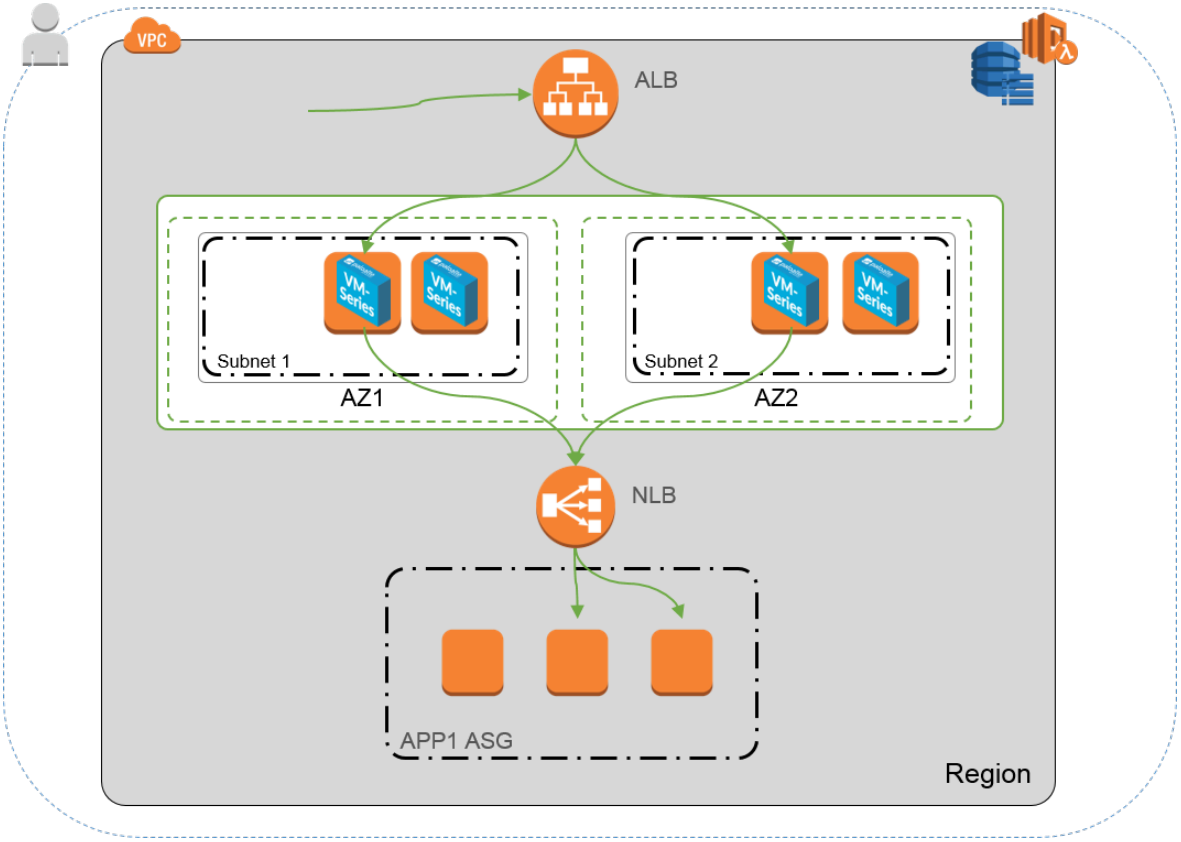

The following application template deploys the network load balancer depicted in the preceding image.

- **Application Template**—The application template deploys a network load balancer (NLB) and one auto scaling group (ASG) with a web server in each AZ.

  The application template is community supported. This template is provided as an example to help you get started with a basic web application. For a production environment, either use your own application template or customize this template to meet your requirements.

These templates allow you to deploy a load balancer sandwich topology with an internet-facing ALB and an internal NLB. The ALB is accessible from the internet and distributes traffic that enters the VPC across a pool of VM-Series firewalls. The firewalls then route traffic using NAT

 ©2024 Palo Alto Networks, Inc.

policy to NLBs, which distributes traffic to an auto scaling tier of web or application servers. The VM-Series firewalls are enabled to publish custom PAN-OS metrics to AWS CloudWatch where you can monitor the health and resource load on the VM-Series firewalls and then use that information to trigger auto scaling events in the appropriate ASGs on firewalls.

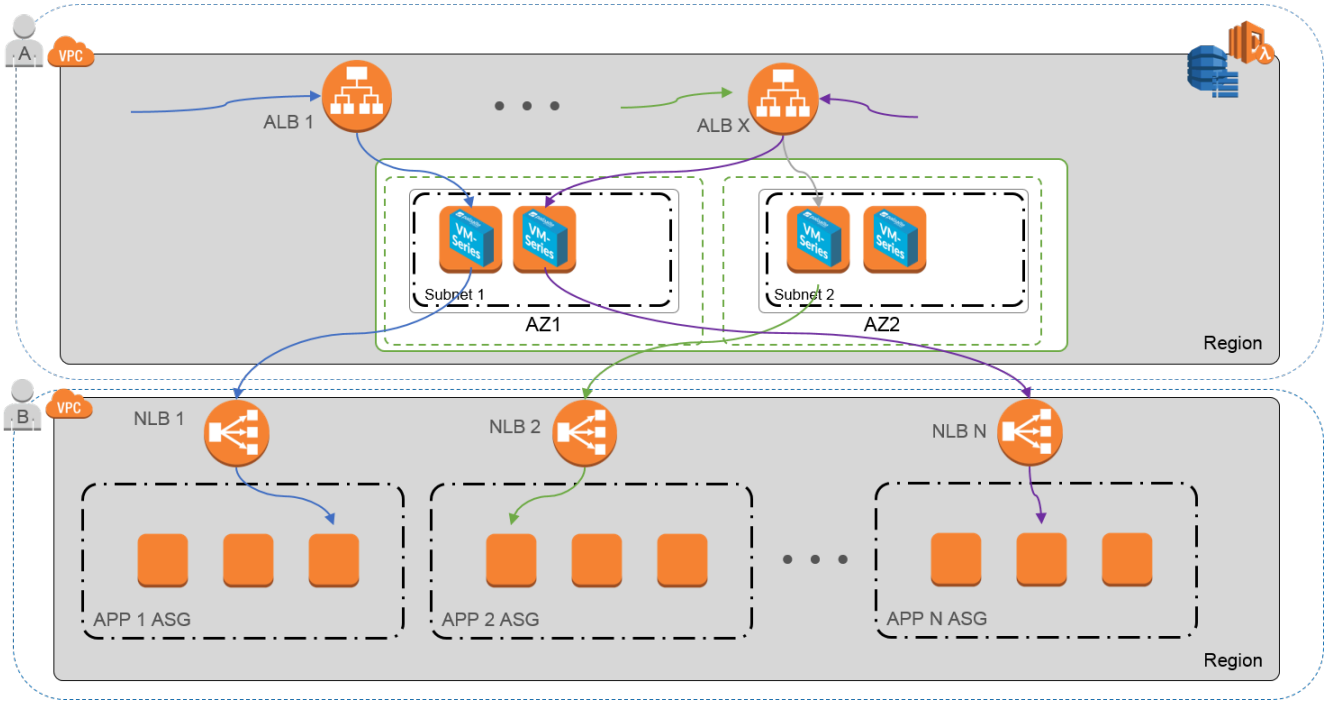

- What Components Does the VM-Series Auto Scaling Template for AWS (v2.0) Leverage?
- How Does the VM-Series Auto Scaling Template for AWS (v2.0) Enable Dynamic Scaling?
- Plan the VM-Series Auto Scaling Template for AWS (v2.0)
- Customize the Firewall Template Before Launch (v2.0)
- Launch the VM-Series Auto Scaling Template for AWS (v2.0)
- Customize the Bootstrap.xml File (v2.0)
- Stack Update with VM-Series Auto Scaling Template for AWS (v2.0)
- Modify Administrative Account and Update Stack

 ©2024 Palo Alto Networks, Inc.

What Components Does the VM-Series Auto Scaling Template for AWS (v2.0) Leverage?

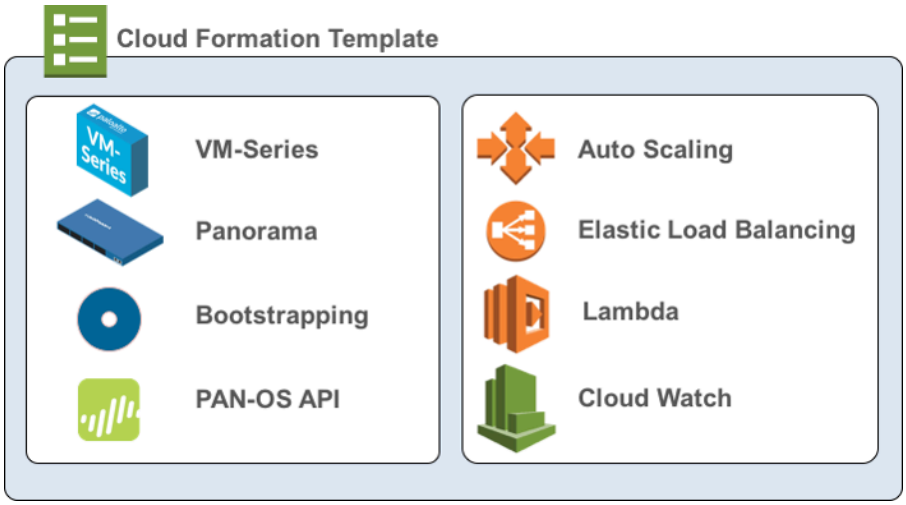

The VM-Series Auto Scaling template for AWS includes the following building blocks:

Building Block	Description
Firewall template (Palo Alto Networks officially supported template)	The `firewall-v2.0.template` deploys a new VPC with subnets, route tables, an AWS NAT gateway, two Availability Zones (AZs), and security groups required for routing traffic across these AZs. This version 2.0 template also deploys an external ALB, and an ASG with a VM-Series firewall in each AZ. Due to the many variations in a production environment that includes but is not limited to a specific number components, such as subnets, availability zones, route tables, and security groups. You must deploy the `firewall-v2.0.template` in a new VPC. *VM-Series Auto Scaling template for AWS does not deploy Panorama and Panorama is optional. Panorama provides ease of policy management and central visibility. If you want to use Panorama to manage the VM-Series firewalls that the solution deploys, you can either use an M-Series appliance or Panorama virtual appliance inside your corporate network or you can use a Panorama virtual appliance on AWS.* This solution includes an AWS NAT gateway that the firewalls use to initiate outbound requests for retrieving updates, connecting to Panorama, and publishing metrics to AWS CloudWatch.
Application template (Community supported template)	The application template deploys an NLB and an ASG with a web server in each AZ. Because the NLB has a unique IP address for each AZ and the NAT policy rule on the firewalls must reference a single IP

 ©2024 Palo Alto Networks, Inc.

Building Block	Description
	address, there is one ASG for each of the two AZs. All firewalls in an ASG use an identical configuration. Version 2.0 of the auto scaling solution includes two application templates: • The `panw_aws_nlb-v2.0.template` allows you to deploy the application template resources within the same VPC as the one in which you deployed the firewall template (same AWS account). • The `panw_aws_nlb_vpcv-2.0.template` allows you to deploy the application template resources in a separate VPC using the same AWS account or multiple AWS accounts.
Lambda functions	AWS Lambda provides robust, event-driven automation without the need for complex orchestration software. In the `firewall-v2.0.template`, AWS Lambda monitors a Simple Queue Service (SQS) to learn about NLBs that publish to the queue. When the Lambda function detects a new NLB, it creates a new NAT policy rule and applies it to the VM-Series firewalls within the ASG. The firewalls have a NAT policy rule for each application and the firewalls use the NAT policy rule (that maps the port to NLB IP address) to forward traffic to the NLB in front of the application web servers. *You need to create the Security policy rule to allow or deny application traffic for your deployment. The sample `bootstrap.xml` file does not include any Security policy rules. You should use Panorama to centrally manage your firewalls and simplify creating Security policy rules.* There are additional functions: • Adds or removes an interface (ENI) when a firewall is launched or terminated. • Deletes all the associated resources when you delete a stack or terminate an instance. • Removes a firewall as a Panorama managed device when there is a scale-in event. • Deactivates the BYOL license when a scale-in event results in a firewall termination. To learn more about the Lambda functions, refer to **http://paloaltonetworks-aws-autoscale-2-0.readthedocs.io/en/latest/**
Bootstrap files The bootstrap.xml file provided in the GitHub repository is provided for testing	This solution requires the init-cfg.txt file and the bootstrap.xml file so that the VM-Series firewall has the basic configuration for handling traffic.

 ©2024 Palo Alto Networks, Inc.

Building Block	Description
and evaluation only. For a production deployment, you must modify the sample credentials in the bootstrap.xml prior to launch.	• The `init-cfg.txt` file includes the mgmt-interface-swap operational command to enable the firewall to receive dataplane traffic on its primary interface (eth0). This auto-scaling solution requires the swapping of the dataplane and management interfaces to enable the ALB to forward web traffic to the auto-scaling tier of VM-Series firewalls. For details, see Management Interface Mapping for Use with Amazon ELB. • The `bootstrap.xml` file enables basic connectivity for the firewall network interfaces and allows the firewall to connect to the AWS CloudWatch namespace that matches the stack name you enter when you launch the template.

To deploy the solution, see Launch the VM-Series Auto Scaling Template for AWS (v2.0).

How Does the VM-Series Auto Scaling Template for AWS (v2.0 and v2.1) Enable Dynamic Scaling?

VM-Series firewall scale in and scale out using VM-Series firewalls that are deployed using auto scaling templates based on custom PAN-OS metrics. The VM-Series firewalls natively publish these metrics to the Amazon CloudWatch console and, based on the metrics you choose for the scaling parameters, you can define CloudWatch alarms and policies to dynamically deploy or terminate instances for managing the application traffic in your AWS deployment.

The firewalls publish metrics to AWS CloudWatch every five minutes (by default). When a monitored metric reaches the configured threshold for the defined time interval, CloudWatch triggers an alarm and initiates an auto-scaling event.

When the auto-scaling event triggers the deployment of a new firewall, the new instance bootstraps at launch and an AWS Lambda function configures the firewall with NAT policy rules. A NAT policy rule is created for each application and the rule references the IP addresses for each network load balancer in your deployment. When the application load balancer receives a request, it forwards the request to the firewall on the assigned TCP port. The firewall then inspects the traffic and forwards it to the corresponding network load balancer, which then forwards the request to a web server in its target group.

Plan the VM-Series Auto Scaling Template for AWS (v2.0 and v2.1)

The items in this checklist are actions and choices you must make to implement this solution.

Planning Checklist for Templates v2.0 and v2.1	
☐ Verify the requirements for deploying the VM-Series Auto Scaling template.	The auto scaling template requires AWS Lambda and S3 Signature versions 2 or 4, and can deploy VM-Series firewalls running supported PAN-OS versions. You need to look up the list of supported regions and the AMI IDs, to provide as an input in the firewall template.

 ©2024 Palo Alto Networks, Inc.

Planning Checklist for Templates v2.0 and v2.1	
❑ Assign the appropriate permissions for the IAM user role.	The user who deploys the VM-Series Auto Scaling template must either have administrative privileges or have the permissions listed in the iam-policy.json to launch this solution successfully. Copy and paste the permissions from this file in to a new IAM policy and then attach the policy to a new or existing IAM role. For a cross-account deployment, to access resources that are in a different AWS accounts, the IAM role for the user who deploys the application template must have full SQS access permissions and a trust relationship that authorizes her to write to the SQS queue that belongs to the firewall template.
❑ Collect the details required for a cross-account deployment.	For a deployment where the firewall template and the application template are in different accounts, the account that hosts the firewall template resources is the trusting account and the other AWS account(s) that hold the application template resources are the trusted accounts. To launch the application template in a cross-account deployment, you need the following information: • Cross-account Role Amazon Resource Name (ARN) of the account in which you are deploying the application template. • External ID, which you defined when creating the IAM role that grants full SQS access to the trusting account. • The 10-digit account number for every AWS account in which you plan to launch the application template. Because the account that hosts the firewall template resources serves as a trusting account, and it owns the resources that the users of the application template need, you need to list the account number for each trusted account that can access the firewall resources.
❑ Create a Support Account on the Palo Alto Networks Support portal, if you don't already have one.	You can opt for the BYOL or PAYG licenses. • For BYOL, you must register an auth code to your Palo Alto Networks support account prior to launching the VM-Series Auto Scaling template and add the auth-code to the `/license` folder with filename as authcodes in the bootstrap package. See Launch the VM-Series Auto Scaling Template for AWS (v2.0) or Launch the Firewall Template (v2.1) for details. • For PAYG, you must register the VM-Series firewalls to activate your support entitlement.
❑ (For PAYG only) Review and accept the End User License	In the AWS Marketplace, search for Palo Alto Networks, and select the bundle you plan to use. The VM-Series firewalls will fail to deploy if you have not accepted the EULA for the bundle you plan to use.

 ©2024 Palo Alto Networks, Inc.

Planning Checklist for Templates v2.0 and v2.1	
Agreement (EULA). Required, if you are launching a VM-Series firewall in an AWS account for the first time.	• Search for **VM-Series Next Generation Firewall Bundle 2, for example**. 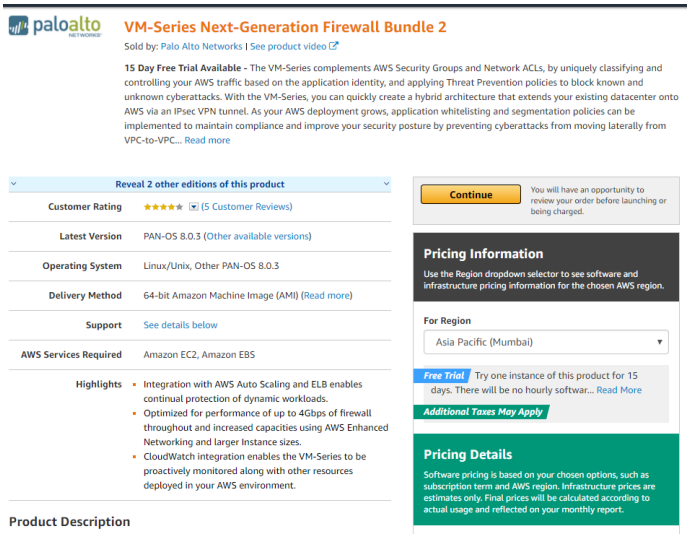 • Click **Continue**, and select **Manual Launch**. Review the agreement and click **Accept Software Terms** to accept the EULA. 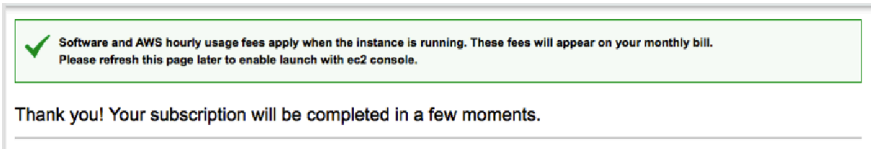 You can now close the browser.
❑ Decide whether you plan to use the public S3 buckets or your private S3 bucket for AWS Lambda, Python scripts, and templates.	Palo Alto Networks provides public S3 buckets in all AWS regions included in the supported regions list. These S3 buckets include all the templates, AWS Lambda code, and the bootstrap files that you need. *Palo Alto Networks recommends using the bootstrap files in the public S3 bucket only for evaluating this solution. For a production deployment, you must create a private S3 bucket for the bootstrap package.* The naming convention for the S3 bucket is `panw-aws-autoscale-v20-<region_name>`. For example, the bucket in the AWS Oregon region is panw-aws-autoscale-v20-us-west-2. To use your private S3 bucket, you must download and copy the templates, AWS Lambda code, and the bootstrap files to your private S3 bucket. You can place all the required files for both the firewall template and the application template in one S3 bucket or place them in separate S3 buckets.

 ©2024 Palo Alto Networks, Inc.

Planning Checklist for Templates v2.0 and v2.1	

❑ Download the templates, AWS Lambda code, and the bootstrap files.	• Get the files for deploying the firewall template (application load balancer and the VM-Series firewalls) from the GitHub repository.
	
	Do not mix and match files across VM-Series Auto Scaling template versions.
	
	• Templates and Lambda code:
	  • panw-aws.zip
	  • firewall-v2.X.template
	• Bootstrap files:
	  • init-cfg.txt
	  • bootstrap.xml
	
	The bootstrap.xml file bundled with this solution is designed to help you get started, and is provided for testing and evaluation only. For a production deployment, you must modify the bootstrap.xml prior to launch.
	• iam-policy: The user who deploys the VM-Series Auto Scaling template must have either the administrative privileges or the permissions listed in this file to successfully launch this solution.
	
	The firewall template is supported by Palo Alto Networks Technical Support.
	• Get the files for deploying the NLB and the web servers from the GitHub repository versions 2.0 or 2.1.
	• Templates:
	  • pan_aws_nlb-2.X.template—Use this template to deploy the application template resources within same VPC as the one in which you deployed the firewall template (same AWS account).
	  • pan_aws_nlb_vpc-2.X.template—Use this template to deploy the application template resources in a different VPC. This template allows you to deploy the resources within the same AWS account or in a different AWS account as long as you have the appropriate permissions to support a cross-account deployment.
	  • pan_nlb_lambda.template—Creates an AWS Network Load balancer, which multiplexes traffic to register scaled-out backend web servers.
	• Lambda code and Python scripts.

❑ Customize the bootstrap.xml	To ensure that your production environment is secure, you must customize the bootstrap.xml file with a unique administrative

 ©2024 Palo Alto Networks, Inc.

Planning Checklist for Templates v2.0 and v2.1	
file for your production environment.	username and password for production deployments. The default username and password are pandemo/demopassword. You can also use this opportunity to create an optimal firewall configuration with interfaces, zones, and security policy rules that meet your application security needs.
❑ Decide whether you want to use Panorama for centralized logging, reporting, and firewall management.	Panorama is an option for administrative ease and is the best practice for managing the firewalls. It is not required to manage the auto scaling tier of VM-Series firewalls deployed in this solution. If you want to use Panorama, you can either a Panorama virtual appliance on AWS or use an M-Series appliance or a Panorama virtual appliance inside your corporate network. *The Panorama must be in Panorama mode and not Management Only mode.* To successfully register the firewalls with Panorama, you must collect the following details: • API key for Panorama—So that AWS Lambda can make API requests to Panorama, you must provide an API key when you launch the VM-Series Auto Scaling template. As a best practice, in a production deployment, create a separate administrative account just for the API call and generate an associated API key. • Panorama IP address—You must include the IP address in the configuration (init-cfg.txt) file. The firewalls must be able to access this IP address from the VPC; to ensure a secure connection, use a direct connect link or an IPSec tunnel. • VM auth key—Allows Panorama to authenticate the firewalls so that it can add each firewall as a managed device. You must include this key in the configuration (init-cfg.txt) file. The vm auth key is required for the lifetime of the deployment. Without a valid key in the connection request, the VM-Series firewall will be unable to register with Panorama. For details on the key, see Generate the VM Auth Key on Panorama. • Template stack name and the device group name to which to assign the firewalls—You must first add a template and assign it to a template stack, create a device group on Panorama, and then include the template stack name and the device group name in the configuration (init-cfg.txt) file.

 ©2024 Palo Alto Networks, Inc.

Planning Checklist for Templates v2.0 and v2.1	
	In order to reduce the cost and scale limits of using Elastic IP addresses, the firewalls do not have public IPs. If you are not using Panorama to manage the firewalls, you must deploy a jump server (a bastion host with an EIP address) that attaches to the Untrust subnet within the VPC to enable SSH and/or HTTPS access to the VM-Series firewalls. By default, this solution includes an AWS NAT gateway that the firewalls use to initiate outbound requests for retrieving updates, connecting to Panorama, and publishing metrics to AWS CloudWatch.
Get started	Launch the VM-Series Auto Scaling Template for AWS (v2.0).

Customize the Firewall Template Before Launch (v2.0 and v2.1)

To simplify the deployment workflow, the firewall displays a limited set of parameters for which you need to provide inputs when launching the template. If you would like to view and customize other options included in the template, you can use a text editing tool such as Notepad or Visual Studio Code to specify values that you prefer before you launch the VM-Series Auto Scaling template for AWS v2.0 or 2.1.

Use the following table to view the list of parameters that you are allowed to customize for your deployment of the auto scaling firewall template for AWS. Modifying parameters from this list is within the official support policy of Palo Alto Networks through the support options that you've purchased.

Parameter	Description	Default Value
CIDR Block for the VPC	The IP address space that you want to use for the VPC. *The subnets you modify below must belong to this VPC CIDR block and be unique.*	192.168.0.0/16
Management Subnet CIDR Block	Comma-delimited list of CIDR blocks for the management subnet of the firewalls.	192.168.0.0/24, 192.168.10.0/24
Untrust Subnet CIDR Block	Comma-delimited list of CIDR blocks for the Untrust subnet.	192.168.1.0/24, 192.168.11.0/24
Trust Subnet CIDR Block	Comma-delimited list of CIDR blocks for the Trust subnet.	192.168.2.0/24, 192.168.12.0/24

 ©2024 Palo Alto Networks, Inc.

Parameter	Description	Default Value
NAT Gateway Subnet CIDR Block	Comma-delimited list of CIDR blocks for the AWS NAT Gateway.	192.168.100.0/24, 192.168.101.0/24
Lambda Subnet CIDR Block	Comma-delimited list of CIDR blocks for the Lambda functions.	192.168.200.0/24, 192.168.201.0/24
Firewall Instance size	AWS Instance Types and size that you want for the VM-Series firewalls in your deployment.	M4.xlarge
Choose your Scaling Parameter *You do not need to modify the template for the scaling parameter. You can set AWS CloudWatch alarms on the AWS console for one or more custom PAN-OS metrics on which you want to trigger autoscaling.*	The template publishes all the following metrics to AWS CloudWatch: • CPU—DataPlane CPU Utilization • AS—Active Sessions • SU—Session Utilization • SSPU—SSL Proxy Utilization • GPU—GlobalProtect Gateway Utilization • GPAT—GlobalProtect Gateway Utilization ActiveTunnels • DPB—Dataplane Packet Buffer Utilization	Dataplane CPU Utilization
Choose time in seconds for Scaling Period	The period in seconds over which the average statistic is applied. Must be a multiple of 60.	900
Maximum VM-Series Instances	Maximum number of VM-Series firewalls in the auto scaling group.	3
Minimum VM-Series Instances	Minimum number of VM-Series firewalls in the auto scaling group.	1
ScaleDown threshold value in percentage/value	Value at which a scale in event is triggered.	20

 ©2024 Palo Alto Networks, Inc.

Parameter	Description	Default Value
ScaleUp threshold value in percentage/value	Value at which scale out event is triggered.	80

Launch the VM-Series Auto Scaling Template for AWS (v2.0)

You can choose to deploy the firewall template in one VPC and the sample application template in the same VPC as the one in which you deployed the firewalls, or in a different VPC.

If the applications that you want to secure belong to a separate AWS account, the sample application template includes support for cross-account deployments. The solution supports a hub and spoke architecture whereby you can deploy the firewall template in one AWS account and use it as a hub to secure your applications (spokes) that belong to the same or to different AWS accounts.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

- Launch the VM-Series Firewall Template
- Launch the Application Template
- (Required only if you deploy more than one internal load balancer) Enable Traffic to the ELB Service (v2.0 and v2.1)

Launch the VM-Series Firewall Template

This workflow tells you how to deploy the application load balancer and the VM-Series firewalls using the firewall template.

This firewall template includes an AWS NAT gateway that the firewalls use to initiate outbound requests for retrieving updates, connecting to Panorama, and publishing metrics to AWS CloudWatch. If you are not using Panorama to manage the firewalls, you must deploy a jump server (a bastion host with an EIP address) that attaches to the Untrust subnet within the VPC to enable SSH and/or HTTPS access to the VM-Series firewalls. This jump server is required because the management interface on the VM-Series firewalls has a private IP address only.

STEP 1 | Review the checklist for Plan the VM-Series Auto Scaling Template for AWS (v2.0).

Make sure that you have completed the following tasks:

- (For PAYG only) Reviewed and accepted the EULA for the PAYG bundle you plan to use.
- (For BYOL only) Obtained the auth code. You need to enter this auth code in the /license folder of the bootstrap package.
- Downloaded the files required to launch the VM-Series Auto Scaling template from the GitHub repository.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Modify the init-cfg.txt file. You must add the device certificate auto-registration PIN to the init-cfg.txt file to automatically install a device certificate when your VM-Series firewall instance is deployed.

vm-series-auto-registration-id=

vm-series-auto-registration-pin-value=

For more details read about the bootstrapping process and the init-cfg.txt file.

If you're using Panorama to manage the firewalls, complete the following tasks:

1. Generate the VM Auth Key on Panorama. The firewalls must include a valid key in the connection request to Panorama. Set the lifetime for the key to 8760 hours (1 year).
2. Open the init-cfg.txt file with a text editor, such as Notepad. Make sure that you do not alter the format as this causes a failure in deploying the VM-Series Auto Scaling template. Add the following information as name-value pairs:
   - IP addresses for the primary Panorama and optionally a secondary Panorama. Enter:

     panorama-server=

     panorama-server-2=
   - Specify the template stack name and the device group to which you want to assign the firewall. Enter:

     tplname=

     dgname=
   - VM auth key. Enter:

     vm-auth-key=
3. Verify that you have not deleted the command for swapping the management interface (mgmt) and the dataplane interface (ethernet 1/1) on the VM-Series firewall on AWS. For example, the file must include name-value pairs as shown here:

   ```
   op-command-modes=mgmt-interface-swap
   vm-auth-key=755036225328715
   panorama-server=10.5.107.20
   panorama-server-2=10.5.107.21
   tplname=FINANCE_TG4
   dgname=finance_dg
   ```
4. Save and close the file.

STEP 3 | (For BYOL only) Add the license auth code in the /license folder of the bootstrap package. For more information see Prepare the Bootstrap Package.

1. Create a new .txt file with a text editor, such as Notepad.
2. Add the authcode for your BYOL licenses to this file,then save the file with authcodes (no file extension) and upload it to the /license folder. The auth code must support the number of firewalls that may be required for your deployment. You must use an auth code bundle instead of individual auth codes so that the firewall can simultaneously

 ©2024 Palo Alto Networks, Inc.

fetch all license keys associated with a firewall. If you use individual auth codes instead of a bundle, the firewall retrieves only the license key for the first auth code included in the file.

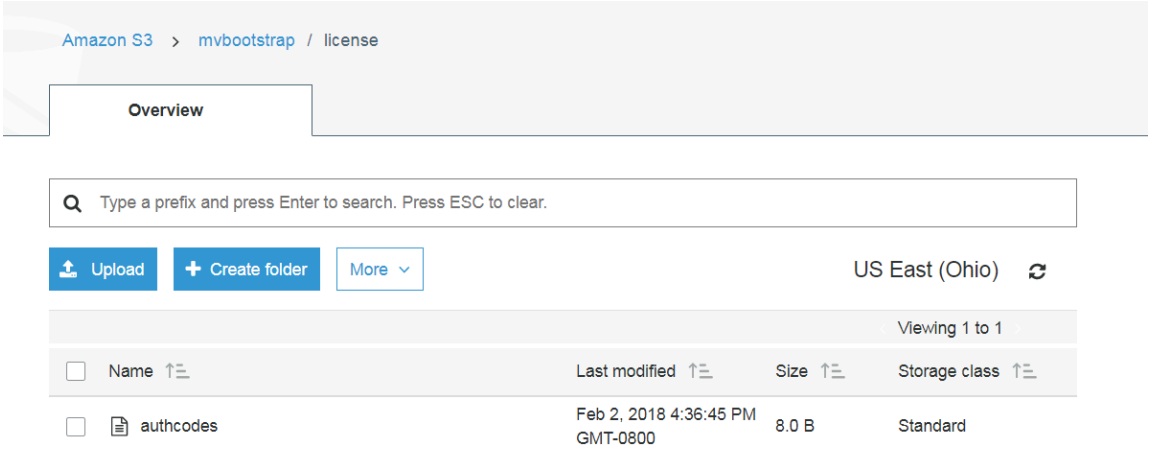

STEP 4 | Change the default credentials for the VM-Series firewall administrator account defined in the bootstrap.xml file.

Required for using the VM-Series Auto Scaling template in a production environment.

The bootstrap.xml file in the GitHub repository is provided for testing and evaluation only. For a production deployment, you must Customize the Bootstrap.xml File (v2.0) prior to launch.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Prepare the Amazon Simple Storage (S3) buckets for launching the VM-Series Auto Scaling template to a production environment.

Make sure to create the S3 buckets in the same region in which you plan to deploy the template; the bootstrapping files hosted in the public S3 bucket are provided only to make it easier for you to evaluate the template.

1. Create a new S3 bucket for the bootstrap files.
   1. Sign in to the AWS Management Console and open the S3 console.
   2. Click **Create Bucket**.
   3. Enter a **Bucket Name** and a **Region**, and click **Create**. The bucket must be at the S3 root level. If you nest the bucket, bootstrapping fails because you cannot specify a path to the location of the bootstrap files.
2. Upload the bootstrap files to the S3 bucket. The bootstrap folders must be in the root folder of the S3 bucket.
   1. Click the name of bucket and then click **Create folder**.
   2. Create the following folder structure for bootstrapping.

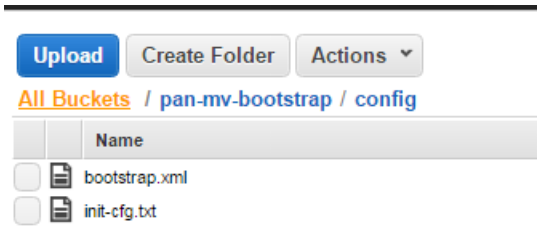

   3. Click the link to open the **config** folder.
   4. Select **Actions** > **Upload** and **Add Files**, browse to select the init-cfg.txt file and bootstrap.xml file, and click **Open**.
   5. Click **Start Upload** to add the files to the config folder. The folder can contain only two files: init-cfg.txt and the bootstrap.xml.

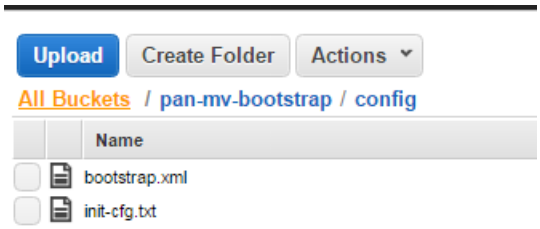

   6. (For BYOL only) Click the link to open the **license** folder and upload the txt file with the auth code required for licensing the VM-Series firewalls.

 ©2024 Palo Alto Networks, Inc.

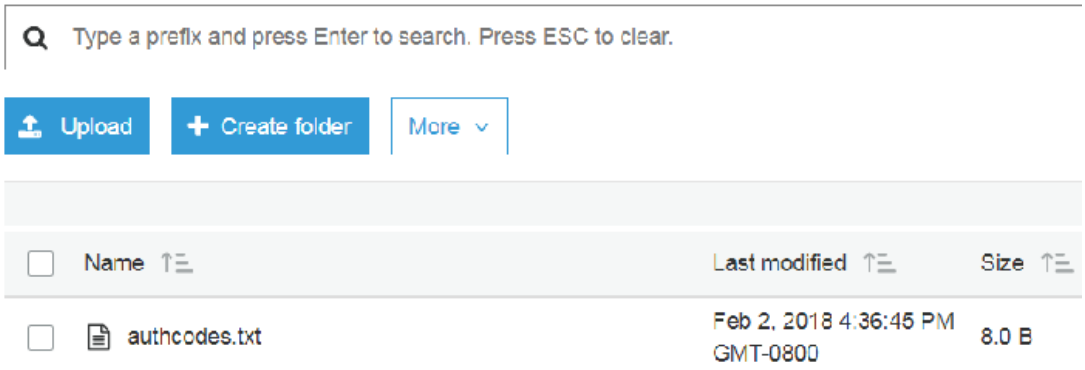

3. Upload the AWS Lambda code (panw-aws.zip file) to an S3 bucket. In this example, the AWS Lambda code is in the same S3 bucket as the bootstrap package.
   1. Click the bucket name.
   2. Click **Add Files** to select the panw-aws.zip file, click **Open**.
   3. Click **Start Upload** to add the zip file to the S3 bucket.

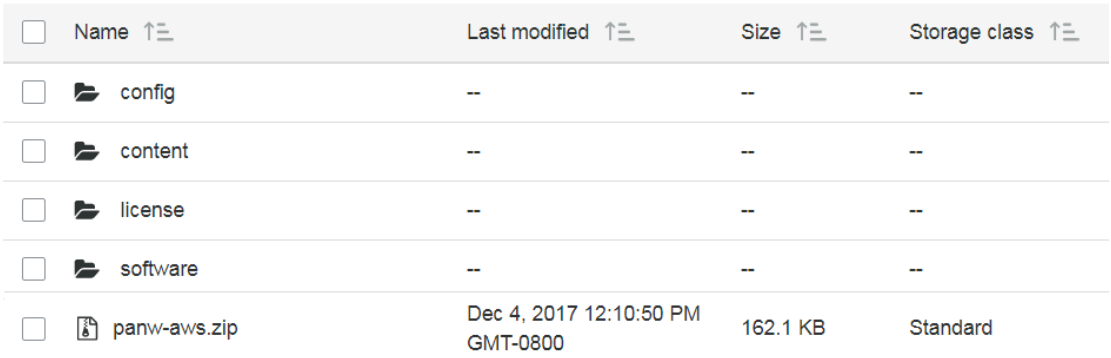

STEP 6 | Select the firewall template.

If you need to Customize the Firewall Template Before Launch (v2.0), do that now and select the modified template.

1. In the AWS Management Console, select **CloudFormation** > **Create Stack**.
2. Select **Upload a template to Amazon S3**, choose the firewall-v2.0.template and click **Open** and **Next**.
3. Specify the **Stack name**. The stack name allows you to uniquely identify all the resources that this template deploys.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Configure the parameters for the VPC.

1. Enter the parameters for the **VPC Configuration** as follows:
   1. Enter a **VPCName**.
   2. Select the two Availability Zones that your setup spans in **Select two AZs**.

STEP 8 | Select your preferences for the VM-Series firewalls.

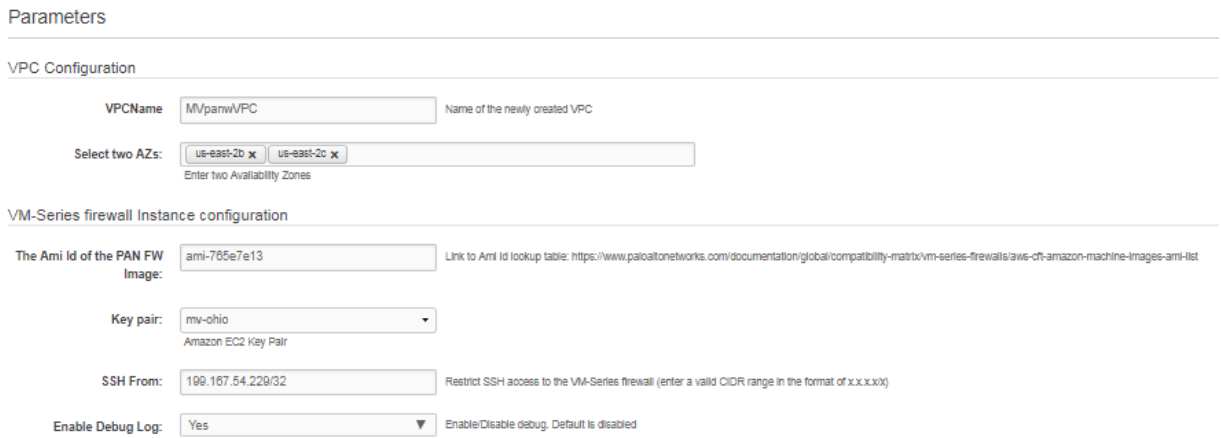

1. Look up the AMI ID for the VM-Series firewall and enter it. Make sure that the AMI ID matches the AWS region, PAN-OS version and the BYOL or PAYG licensing option you opted to use.
2. Select the EC2 **Key pair** (from the drop-down) for launching the firewall. To log in to the firewalls, you must provide the name of this key pair and the private key associated with it.
3. Restrict SSH access to the firewall's management interface. Make sure to supply a CIDR block that corresponds to your dedicated management IP addresses or network. Do not make the allowed source network range larger than necessary and do not ever configure the allowed source as 0.0.0.0/0. Verify your IP address before configuring it on the template to make sure that you do not lock yourself out.
4. Select **Yes** if you want to **Enable Debug Log**. Enabling the debug log generates more verbose logs that help with troubleshooting issues with the deployment. These logs are generated using the stack name and are saved in AWS CloudWatch.

By default, the template uses CPU utilization as the scaling parameter for the VM-Series firewalls. Custom PAN-OS metrics are automatically published to the CloudWatch namespace that matches the stack name you specified earlier.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Specify the name of the Amazon S3 bucket(s).

You can use one S3 bucket for the bootstrap package and the zip file.

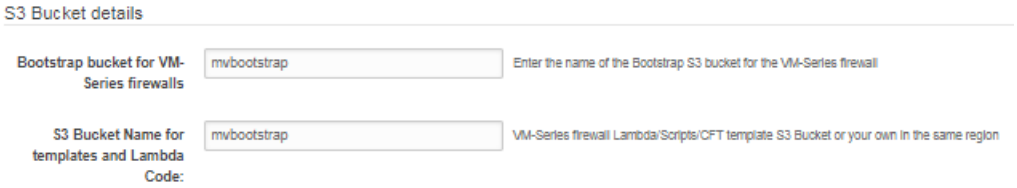

1. Enter the name of the S3 bucket that contains the bootstrap package.

   If the bootstrap bucket is not set up properly or if you enter the bucket name incorrectly, the bootstrap process fails and you cannot be able to log in to the firewall. Health checks for the load balancers also fail.
2. Enter the name of the S3 bucket that contains the panw-aws.zip file.

STEP 10 | Specify the keys for enabling API access to the firewall and Panorama.

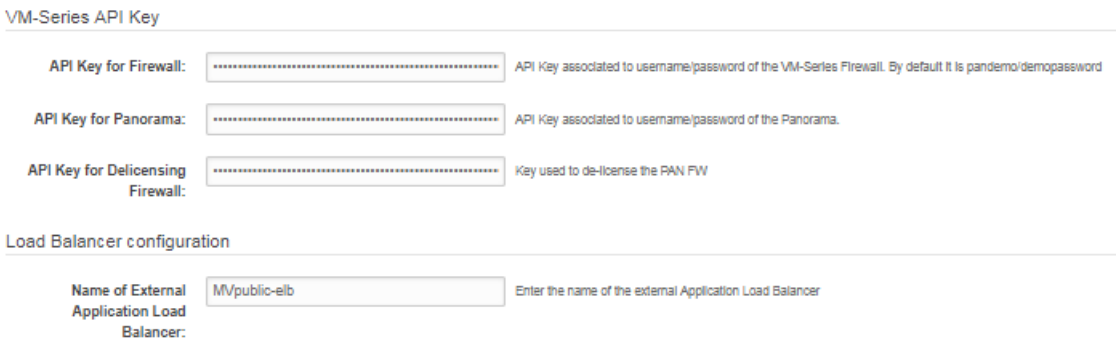

1. Enter the key that the firewall must use to authenticate API calls. The default key is based on the sample bootstrap.xml file and you should only use it for testing and evaluation. For a production deployment, you must create a separate PAN-OS login just for the API call and generate an associated key.
2. Enter the API Key to allow AWS Lambda to make API calls to Panorama, if you are using Panorama for centralized management. For a production deployment, you should create a separate login just for the API call and generate an associated key.
3. Copy and paste the license deactivation API key for your account. This key is required to successfully deactivate licenses on your firewalls when a scale-in event occurs. To get this key:
   1. Log in to the Customer Support Portal.
   2. Select **Assets** > **API Key Management**.
   3. Copy the API key.

STEP 11 | Enter the name for the application load balancer.

STEP 12 | (Optional) Apply tags to identify the resources associated with the VM-Series Auto Scaling template.

Add a name-value pair to identify and categorize the resources in this stack.

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Review the template settings and launch the template.

1. Select **I acknowledge that this template might cause AWS CloudFormation to create IAM resources**.
2. Click **Create** to launch the template. The CREATE_IN_PROGRESS event displays.
3. On successful deployment the status updates to CREATE_COMPLETE.

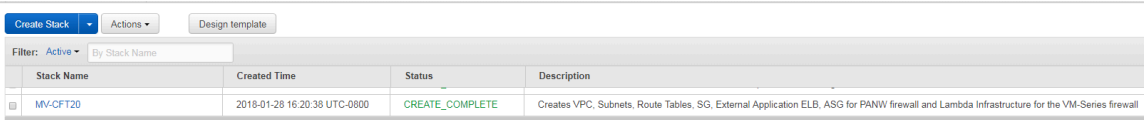

Unless you customized the template, the VM-Series Auto Scaling template launches an ASG that includes one VM-Series firewall in each AZ, behind the application load balancer.

 ©2024 Palo Alto Networks, Inc.

STEP 14 | Verify that the template has launched all required resources.

1. On the AWS Management Console, select the stack name to view the **Output** for the list of resources.

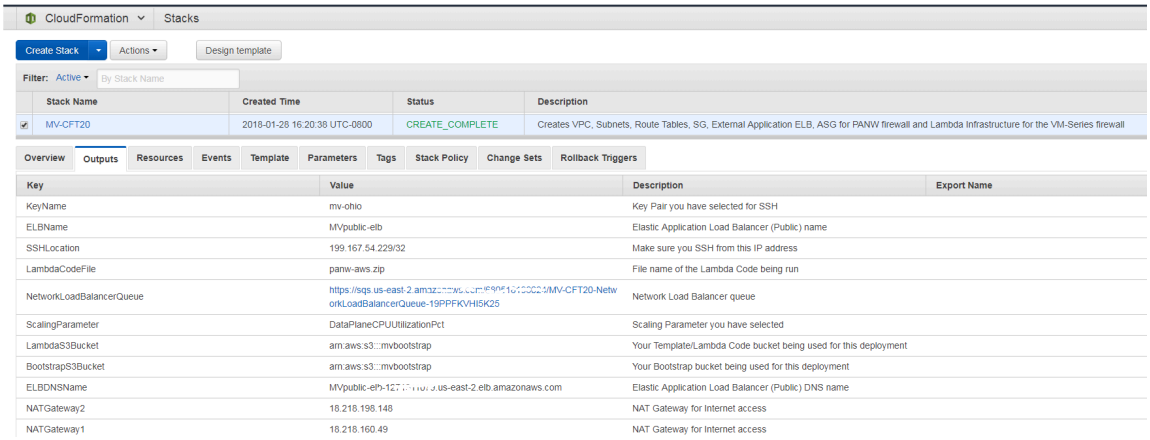

2. On the EC2 Dashboard, select **Auto Scaling Groups**. Verify that in each AZ, you have one ASG for the VM-Series firewalls with the one firewall in each ASG. The ASG name prefix includes the stack name.
3. Log in to the VM-Series firewall. You must deploy a jump server or use Panorama to access the web interface on the firewall.

- *It can take up to 20 minutes for the firewalls to boot up and be available to handle traffic.*
- *When you finish testing or a production deployment, the only way to ensure charges stop occurring is to completely delete the stack. Shutting down instances, or changing the ASG maximum to 0 is not sufficient.*

 ©2024 Palo Alto Networks, Inc.

STEP 15 | Save the following information. You need to provide these values as inputs when deploying the application template.

- IP addresses of the NAT Gateway in each AZ. You need this IP address to restrict HTTP access to the web servers if you deploy the application in a different VPC. Specifying this IP address ensures that the firewall secures access your applications in a different VPC, and that nobody can bypass the firewall to directly access the web server. The sample application template (panw_aws_nlb_vpc-2.0.template) displays a template validation error if you do not enter the NAT Gateway IP addresses; you must enter the IP addresses as a comma-separated list.
- Network Load Balancer SQS URL. An AWS Lambda function in the firewall stack monitors this queue so that it can learn about any network load balancers that you deploy, and create NAT policy rules (one per application) on the VM-Series firewalls that enable the firewalls to send traffic to the network load balancer IP address.

Launch the Application Template

The application template allows you to complete the sandwich topology and is provided so that you can evaluate the auto scaling solution. This application template deploys a network load balancer and a pair of web servers behind the auto scaling group of VM-Series firewalls, which you deployed using the firewall template. The web servers in this template have a public IP address for direct outbound access to retrieve software updates. Use this template to evaluate the solution, but build your own template to deploy to production. For a custom template, make sure to enable SQS Messaging Between the Application Template and Firewall Template.

When launching the application template, you must select the template based on whether you want to deploy the application template within the same VPC (panw_aws_nlb-2.0.template) in which you deployed the firewall template or in a separate VPC (panw_aws_nlb_vpc-2.0.template). For a separate VPC, the template provides supports for cross-account deployments. A cross-account deployment requires you to create an IAM role and enable permissions and trust relationship between the trusting AWS account and the trusted AWS account, and the account information is required as input when launching the template.

STEP 1 | (Required only for a cross-account deployment) Create the IAM role. Refer to AWS documentation.

This role grants access to a user who belongs to a different AWS account. This user requires permissions to access the Simple Queue Service (SQS) resource in the firewall template. The firewall uses this queue to learn about each network load balancer that you deploy so that

 ©2024 Palo Alto Networks, Inc.

it can create NAT policy to send traffic to the web servers that are behind the network load balancer.

- For **Account ID**, type the AWS account ID of the account into which you are deploying the application template. Specifying that account ID allows you to grant access to the resources in your account that hosts the firewall template resources.
- Select **Require external ID** and enter a value that is a shared secret. Specifying an external ID allows the user to assume the role only if the request includes the correct value.
- Choose **Permissons** to allow **Amazon SQS Full Access**.

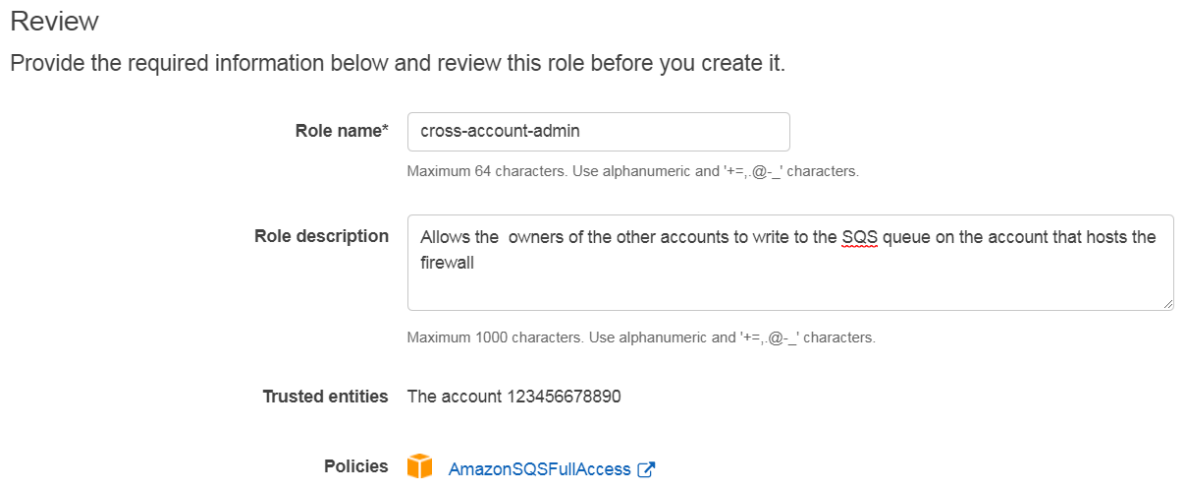

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Use the Palo Alto Networks public S3 bucket or prepare your private (S3) bucket for launching the application template.

1. Create a zip file with all the files in the GitHub repository, excluding the three .template files, named nlb.zip in the screenshot below.
2. Upload the zip file to the S3 bucket you created earlier or to a new bucket.

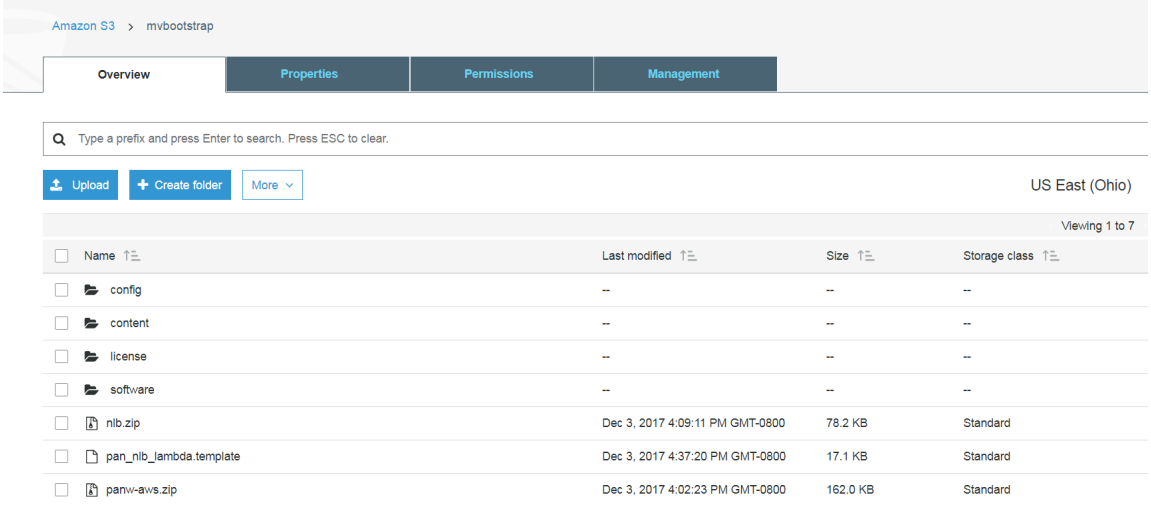

 ©2024 Palo Alto Networks, Inc.

3. Copy the pan_nlb_lambda template into the same bucket to which you copied the nlb.zip file.

STEP 3 | Select the application template to launch.

1. In the AWS Management Console, select **CloudFormation** > **Create Stack**.
2. Select **Upload a template to Amazon S3**, to choose the panw_aws_nlb-2.0.template to deploy the resources that the template launches within the same VPC as the firewalls, or the panw_aws_nlb_vpc-2.0.template to deploy the resources in to a different VPC. Click **Open** and **Next**.
3. Specify the **Stack name**. The stack name allows you to uniquely identify all the resources that are deployed using this template.

STEP 4 | Configure the parameters for the VPC and network load balancer.

1. Select the two Availability Zones that your setup will span in **Select list of AZ**. If you are deploying within the same VPC make sure to select the same Availability Zones that you selected for the firewall template.
2. Enter a **CIDR Block for the VPC**. The default CIDR is 192.168.0.0/16.

3. (Only if you are using the panw_aws_nlb-2.0.template to deploy the applications within the same VPC)

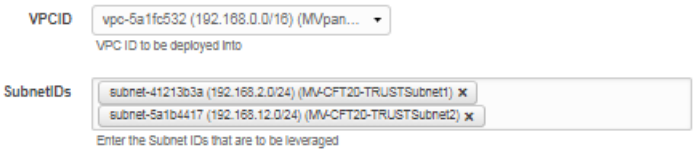

Select the **VPC ID** and the **Subnet IDs** associated with the trust subnet on the firewalls in each AZ. The network load balancer is attached to the trust subnet on the firewalls, to complete the load balancer sandwich topology.

4. Enter a name for the network load balancer.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Configure the parameters for AWS Lambda.

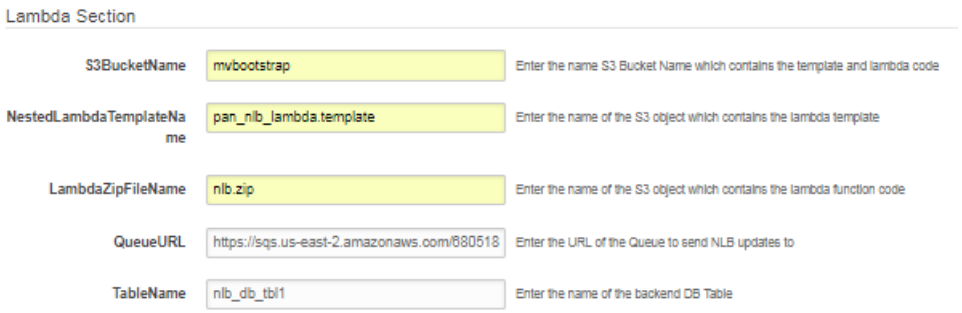

1. Enter the S3 bucket name where nlb.zip and the pan_nlb_lambda.template is stored.
2. Enter the name of the pan_nlb_lambda.template and the zip file name.
3. Paste the SQS URL that you copied earlier.
4. Enter a unique **TableName**. This table stores a mapping of the port and IP address for the applications associated with the network load balancer in your deployment.

   When you delete the application stack this table is deleted. Therefore, if multiple instances of the network load balancer write to the same table and the table is deleted, the NAT rules on the firewalls not function properly and the application traffic maybe be inaccurately forwarded to the wrong port/network load balancer.

STEP 6 | Modify the web server EC2 instance type to meet your deployment needs.

STEP 7 | Select the EC2 **Key pair** (from the drop-down) for launching the web servers. To log in to the web servers, you must provide the key pair name and the private key associated with it.

STEP 8 | (Only if you are using the panw_aws_nlb_vpc-2.0.template) Lock down access to the web servers.

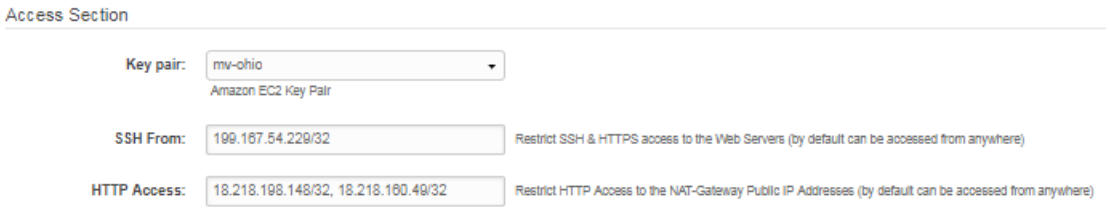

1. Restrict **SSH From** access to the web servers. Only the IP addresses you list here can log in to the web servers.
2. Restrict HTTP access to the web servers. Enter the public IP addresses of the NAT gateway from the firewall template output, and make sure to separate IP addresses with commas. Entering the NAT gateway IP address allows you to ensure that all web traffic to the application servers are secured by the VM-Series firewalls.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | (Only if you are using the panw_aws_nlb_vpc-2.0.template) Configure the other parameters requires to launch the application template stack in a different VPC.

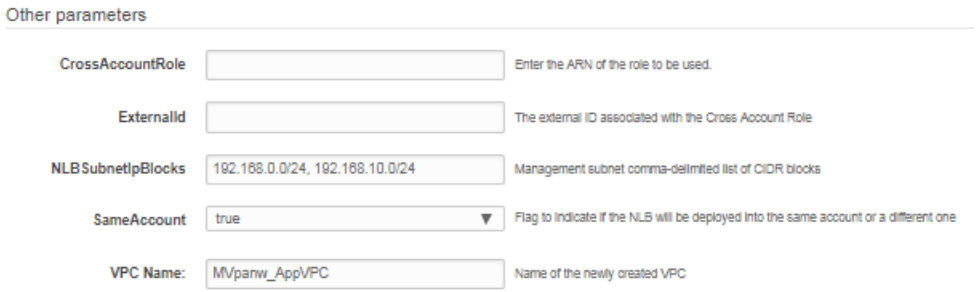

1. Select SameAccount **true** if you are deploying this application template within the same AWS account as the firewall template, and leave the cross account role and external ID blank; select **false** for a cross-account deployment.

   For a cross-account deployment, enter the Amazon Resource Number (ARN) for the **CrossAccountRole** and **ExternalId** that you defined in (Required only for a cross-account deployment) Create the IAM role. Refer to AWS documentation.You can get the ARN from **Support** > **Support Center** on the AWS Management Console.
2. Enter the **VPC Name**in which you want to deploy the application template resources.
3. Optional Change the **NLBSubnetIPBlocks** for the Management subnet for the network load balancer.

STEP 10 | Review the template settings and launch the template.

STEP 11 | Verify that the network load balancer is deployed and in a ready state.

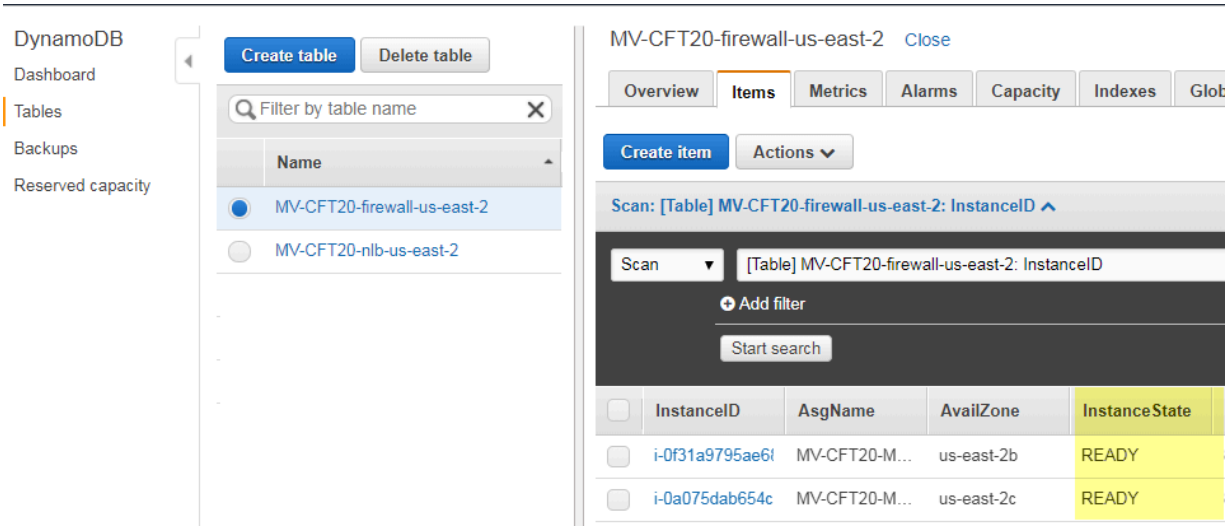

 ©2024 Palo Alto Networks, Inc.

STEP 12 | Get the **DNS name** for the application load balancer, and enter it into a web browser.

For example: http://MVpublic-elb-123456789.us-east-2.elb.amazonaws.com/

When the web page displays, you have successfully launched the auto scaling template.

STEP 13 | Verify that each firewall has a NAT policy rule to the IP address of each network load balancer.

When you deploy the application template to launch another instance of a network load balancer and pair of web servers, the firewall learns about the port allocated for the next network load balancer instance and creates another NAT policy rule. So, if you deploy the application template three times, the firewall has three NAT policy rules for ports 81, 82, and 83.

STEP 14 | If you have launched the application template more than once, you need to Enable Traffic to the ELB Service.

Enable Traffic to the ELB Service (v2.0 and v2.1)

If you add a second or additional internal load balancers (ILBs) in your deployment, you must complete additional configuration so that the internal load balancer, the VM-Series firewalls auto scaling groups, and the web servers can report as healthy and traffic is load balanced across all your AWS resources.

In v2.0, the ILB can only be a network load balancer. In v2.1 the ILB can be an application load balancer or a network load balancer.

STEP 1 | On the AWS management console, verify the ports allocated for each network balancer on the DynamoDB table.

When you launch a new internal load balancer, the application template must send an SQS message to the SQS URL you provided as input when you launched the template. The AWS Lambda function in the firewall template monitors the SQS and adds the port mapping to the DynamoDB table for the firewall template. Starting at port 81, the port allocated for every

 ©2024 Palo Alto Networks, Inc.

additional internal load balancer you deploy increments by 1. So, the second internal load balancer uses port 82, and the third port uses port 83.

1. Select the **DynamoDB** service on the AWS management console.
2. Select **Tables** and click the table that matches the stack name for your firewall template. For example, MV-CFT20-firewall-us-east-2.

   In the Items list, view the ports used by the internal load balancers that are publishing to the SQS associated with the firewall template.

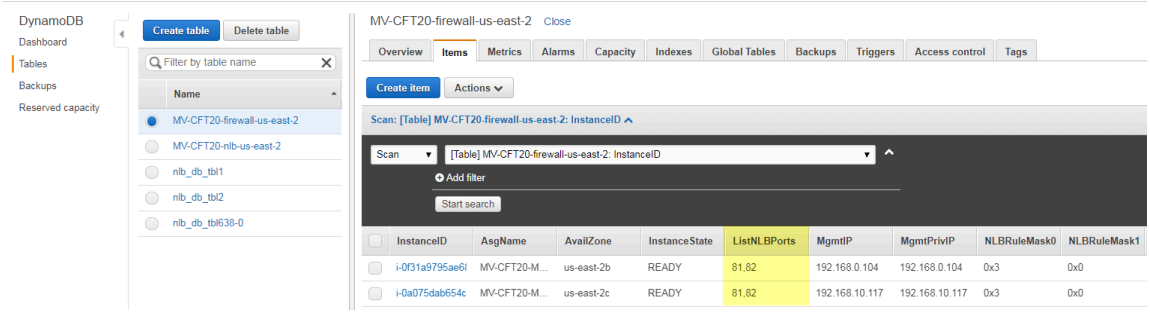

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Create a target group. The internal load balancer sends requests to registered targets using the port and protocol that you specify for the servers in the target group.

When you add a new target group, use the port information that you verified on the DynamoDB table.

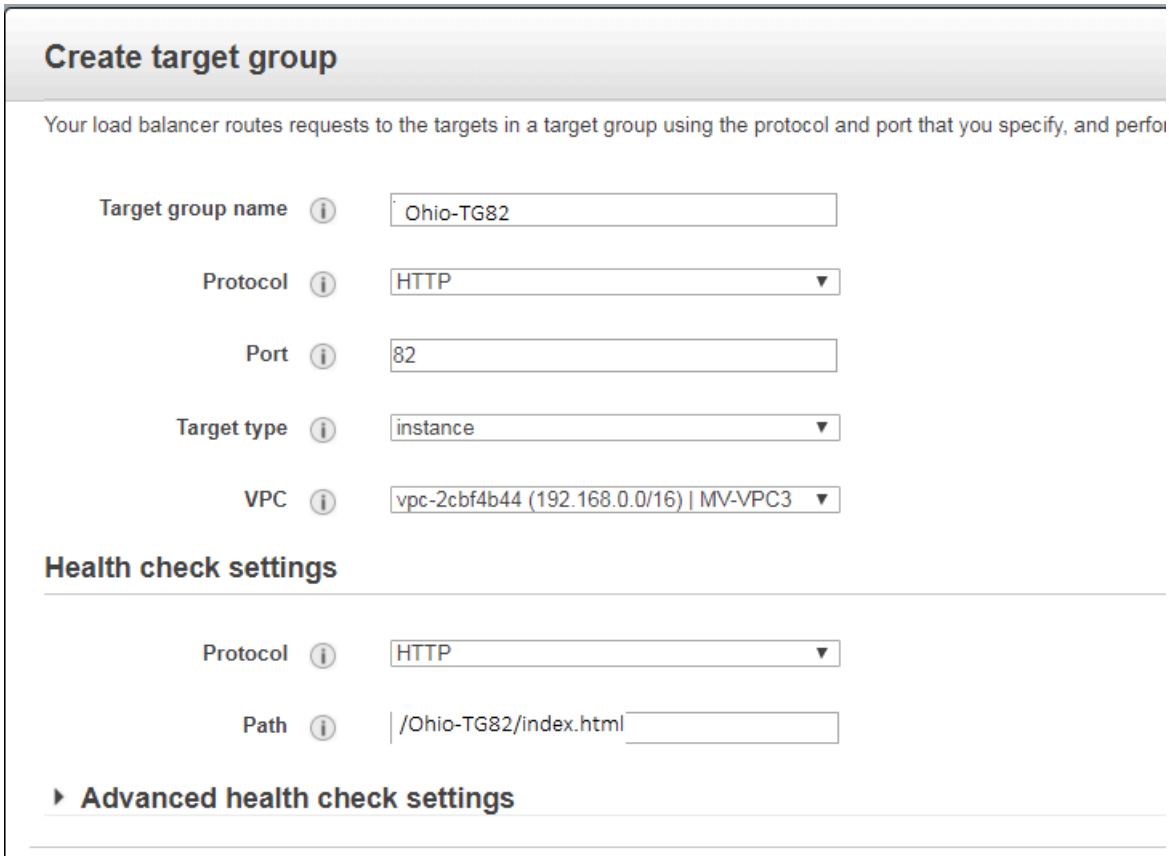

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Edit the listener rules on the internal load balancer to route requests to the target web servers.

1. On the AWS management console, select **Load Balancers** in the Load Balancing section, and select the internal load balancer that matches your stack name.
2. Select **View/edit rules** to modify the rules for the listener.
3. Select **Insert rule** and add a path-based route to forward traffic to the target group you defined above as follows:

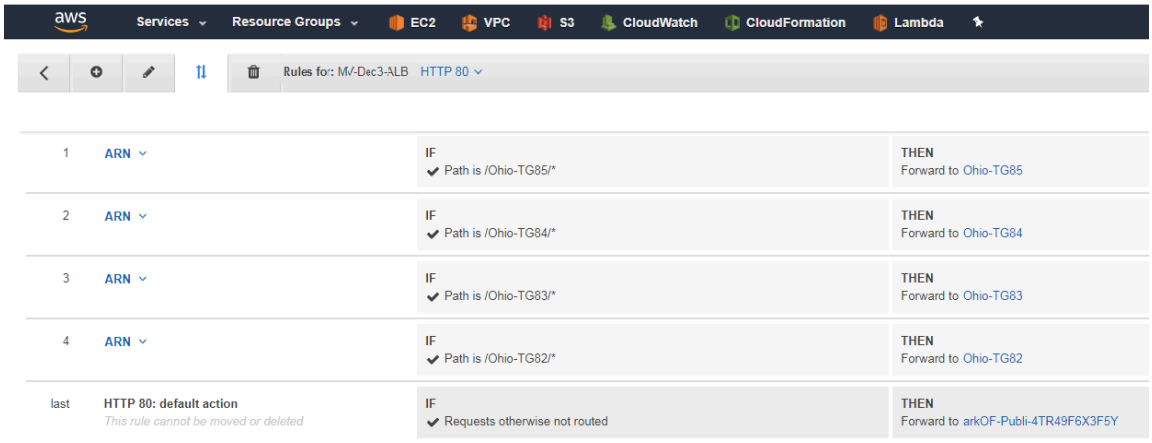

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Attach the target group to both VM-Series firewalls auto scaling groups.

1. Select **Auto Scaling Groups** in the Auto Scaling section and select an auto scaling group that matches the stack name.
2. Select **Details** > **Edit**and select the new target group from the **Target Groups** drop-down.

STEP 5 | Log in to each web server that was deployed by the application template, create a new directory with the target group name and copy the index.html file into the directory. Until you set up the path to the index.html file, the health check for this web server reports as unhealthy.

```
sudo su
cd/var/www/html
mkdir <target-groupname>
cp index.html <target-groupname>
```

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Verify the health status of the web servers.

Select **Auto Scaling Groups**, and use the application stack name to find the webserver auto scaling group to verify that the web servers are reporting healthy.

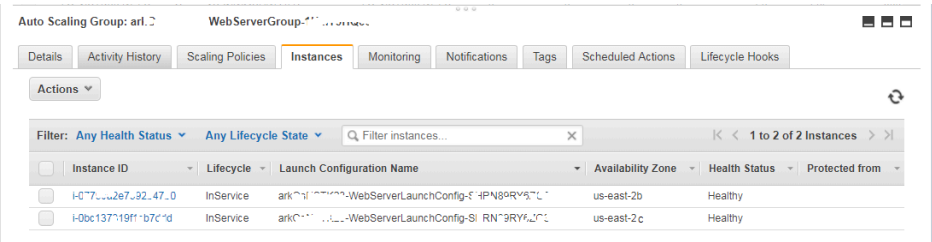

Customize the Bootstrap.xml File (v2.0)

The bootstrap.xml file provided in the GitHub repository uses a default username and password for the firewall administrator. Before deploying the VM-Series Auto Scaling template in a production environment, at a minimum, you must create a unique username and password for the administrative account on the VM-Series firewall. Optionally, you can fully configure the firewall with zones, policy rules, security profiles and export a golden configuration snapshot. You can then use this configuration snapshot as the bootstrap.xml file for your production environment.

You have two ways to customize the bootstrap.xml file for use in a production environment:

- **Option 1**: Launch a VM-Series firewall on AWS using the bootstrap files provided in the GitHub repository, modify the firewall configuration and export the configuration to create a new bootstrap.xml file for the VM-Series Auto Scaling template. See Use the GitHub Bootstrap Files as Seed.
- **Option 2**: Launch a new VM-Series firewall on AWS without using the bootstrap files, add a NAT policy rule to ensure that the VM-Series firewall handles traffic properly, and export the configuration to create a new bootstrap.xml file for the VM-Series Auto Scaling template. See Create a new Bootstrap File from Scratch.

If you have deployed the template and now need to change the credentials for the administrative user or add a new administrative user and update the template stack, see Modify Administrative Account and Update Stack.

Create a new Bootstrap File from Scratch

Launch a new VM-Series firewall on AWS using the AMI for a supported PAN-OS version (see the compatibility matrix for Panorama plugins), without using the sample bootstrap.xml file, and export the configuration to create a new bootstrap.xml file for use with the VM-Series Auto Scaling template v2.0.

STEP 1 | Deploy the VM-Series Firewall on AWS (no bootstrapping required) and use the public IP address to SSH into the Command Line Interface (CLI) of the VM-Series firewall. You will need to configure a new administrative password for the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Log in to the firewall web interface.

STEP 3 | (Optional) Configure the firewall. You can configure the dataplane interfaces, zones and policy rules.

STEP 4 | **Commit** the changes on the firewall.

STEP 5 | Export the configuration file and name it as **`bootstrap.xml`**. (**Device** > **Setup** > **Operation** > **Export Named Configuration Snapshot**).

STEP 6 | Download the bootstrap.xml file from the GitHub repository, open it with a text editing tool, and copy lines 353 to 356. These lines define the AWS CloudWatch namespace to which the firewall publishes custom PAN-OS metrics that are required for the firewalls to auto scale.

STEP 7 | Edit the configuration file you exported earlier to include the AWS CloudWatch information.

Search for `</management>` and paste the lines 353 to 356 after `</management>`.

```
352          </management>
353          <aws-cloudwatch>
354            <enabled>yes</enabled>
355            <name>autoscale-default-panw-asg-name</name>
356          </aws-cloudwatch>
357        </setting>
```

STEP 8 | Delete the management interface configuration.

1. Search for `</service>` and delete the ip-address, netmask and default gateway that follow.
2. Search for `</type>` and delete the ip-address, netmask, default gateway, and public-key that follow.

```
326          </service>
327          <ip-address>192.168.10.16</ip-address>
328          <netmask>255.255.255.0</netmask>
329          <default-gateway>192.168.10.1</default-gateway>
330          <hostname>PA-VM</hostname>
331        </system>
332        <setting>
333          <config>
334            <rematch>yes</rematch>
335          </config>
336          <management>
337            <hostname-type-in-syslog>FQDN</hostname-type-in-syslog>
338            <initcfg>
339              <type>
340                <dhcp-client>
341                  <send-hostname>yes</send-hostname>
342                  <send-client-id>no</send-client-id>
343                  <accept-dhcp-hostname>no</accept-dhcp-hostname>
344                  <accept-dhcp-domain>no</accept-dhcp-domain>
345                </dhcp-client>
346              </type>
347              <ip-address>192.168.10.16</ip-address>
348              <netmask>255.255.255.0</netmask>
349              <default-gateway>192.168.10.1</default-gateway>
350              <public-key>c3NoLXJzYSBBQUFBQjNOemFDMXljMkVBQUFBREFRQUJBQUFCQVFDQTRCSjJwZFBSZlh0TjF2SDVqWW5GRUdYTVdvTmZlaU1FcCtBSlZRaVU4c2hEMHBmSUtOVTVSeHdGRFd4OVZZckRRRFVrLzQ5VDdkeThXcXorcXZ1em44dlFpamY1RDI
VEeHB2WXk3ZWtiam1IallYdTVXUFR4MnZSaXdiMmVzcS91K3FXbm9hSlQ1cXdjU2srbHRxN0prVjlGcC9HSy9jQkRDT0FqOVhmSHMvWi8xQ0VZRk9uZ0UrNTdSL2VwSjFFWWtxZXlLczZuRTBvbWxRajBHSmNhb1FMcUlqZDZmQW5Dcm93dEZ4Y3c3YWVTR
351            </initcfg>
352          </management>
```

STEP 9 | Save the file. You can now proceed with Launch the VM-Series Auto Scaling Template for AWS (v2.0).

Use the GitHub Bootstrap Files as Seed

Launch a VM-Series firewall on AWS from the AWS Marketplace using the bootstrap files provided in the GitHub repository, modify the firewall configuration for your production environment. Then, export the configuration to create a new bootstrap.xml file that you can now use for the VM-Series Auto Scaling template.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | To launch the firewall see Bootstrap the VM-Series Firewall on AWS.

STEP 2 | Add an elastic network interface (ENI) and associate an elastic IP address (EIP) to it, so that you can access the web interface on the VM-Series firewall. See Launch the VM-Series Firewall on AWS for details.

STEP 3 | Use the EIP address to log in to the firewall web interface with admin as the username and password.

STEP 4 | Add a secure password for the administrative user account (**Device** > **Local User Database** > **Users**).

STEP 5 | (Optional) Configure the firewall for securing your production environment.

STEP 6 | **Commit** the changes on the firewall.

STEP 7 | Generate a new API key for the administrator account. Copy this new key to a new file. You will need to enter this API key when you launch the VM-Series Auto Scaling template; the AWS services use the API key to deploy the firewall and to publish metrics for auto scaling.

STEP 8 | Export the configuration file and save it as **`bootstrap.xml`**. (**Device** > **Setup** > **Operation** > **Export Named Configuration Snapshot**).

STEP 9 | Open the bootstrap.xml file with a text editing tool and delete the management interface configuration.

```
326          </service>
327          <ip-address>192.168.10.16</ip-address>
328          <netmask>255.255.255.0</netmask>
329          <default-gateway>192.168.10.1</default-gateway>
330          <hostname>PA-VM</hostname>
331        </system>
332        <setting>
333          <config>
334            <rematch>yes</rematch>
335          </config>
336          <management>
337            <hostname-type-in-syslog>FQDN</hostname-type-in-syslog>
338            <initcfg>
339              <type>
340                <dhcp-client>
341                  <send-hostname>yes</send-hostname>
342                  <send-client-id>no</send-client-id>
343                  <accept-dhcp-hostname>no</accept-dhcp-hostname>
344                  <accept-dhcp-domain>no</accept-dhcp-domain>
345                </dhcp-client>
346              </type>
347              <ip-address>192.168.10.16</ip-address>
348              <netmask>255.255.255.0</netmask>
349              <default-gateway>192.168.10.1</default-gateway>
350              <public-key>c3NoLXJzYSBBQUFBQjNOemFDMXljMkVBQUFBREFRQUJBQUFCQVFDQTRCSjJwZFBSZlh0TjF2SDVqWW5GRUdYTVdvTmZlaU1FcCtBSlZRaVU4c2hEMHBmSUtOVTVSeHdGRFd4OVZZckRRRFVrLzQ5VDdkeThXcXorcXZ1em44dlFpamY1RDJ4dldEaHFWQ3
                VEeHB2WXk3ZWtiam1IallYdTVXUFR4MnZSaXdiMmVzcS91K3FXbm9hSlQ1cXdjU2srbHRxN0prVjlGcC9HSy9jQkRDT0FqOVhmSHMvWi8xQ0VZRk9uZ0UrNTdSL2VwSjFFWWtxZXllLczZuRTBvbWxRajBHSmNhb1FMcUlqZDZmQW5Dcm93dEZ4Y3c3YWVTRnhZdDdXUlFY
351            </initcfg>
352          </management>
```

STEP 10 | (Required if you exported a PAN-OS 8.0 configuration) Ensure that the setting to validate the Palo Alto Networks servers is disabled. Look for `<server-verification>no</server-verification>`.

STEP 11 | If the check is `yes`, change it to **`no`**.

STEP 12 | Save the file. You can now proceed with Launch the VM-Series Auto Scaling Template for AWS (v2.0).

SQS Messaging Between the Application Template and Firewall Template

So that the VM-Series firewalls deployed using the firewall-v2.0.template can detect and send traffic to the network load balancers to which you want to automatically distribute incoming

 ©2024 Palo Alto Networks, Inc.

traffic, the firewall template includes a lambda function that monitors a Simple Queue Service for messages. The message allows the lambda function to learn about a new network load balancer and then automatically create a NAT policy rule on the firewall to send traffic to the IP address of the network load balancer. In order to route traffic properly within the AWS infrastructure, the message must also include basic information on the DNS, VPC ID, and the AZ to which the network load balancer belongs.

If you are building your own application template, you must set up your application template to post two types of messages to the SQS URL that the firewall template in the VM-Series autoscaling template version 2.0 uses to learn about network load balancers to which it must distribute traffic in your environment:

- ADD-NLB message that informs the firewalls when a new network load balancer is available.
- DEL-NLB message that informs the firewalls when a network load balancer has been terminated and is no longer available.

The following examples of each message type includes sample values. You need to modify these message with values that match your deployment.

ADD-NLB Message

```
msg_add_nlb= { 'MSG-TYPE': 'ADD-NLB', 'AVAIL-ZONES': [{'NLB-
IP':'192.168.2.101', 'ZONE-NAME':'us-east-2a', 'SUBNET-ID':
 'subnet-2a566243'}, {'NLB-IP':'192.168.12.101', 'ZONE-NAME':'us-
east-2b', 'SUBNET-ID': 'subnet-2a566243 '}], 'DNS-NAME':
 'publicelb1-2119989486.us-east-2.elb.amazonaws.com', 'VPC-ID':
 'vpc-42ba9f2b', 'NLB-NAME': 'publicelb1' }
```

DEL-NLB Message

```
msg_del_nlb= { 'MSG-TYPE': 'DEL-NLB', 'DNS-NAME':
 'publicelb1-2119989486.us-east-2.elb.amazonaws.com', }
```

Refer to the AWS documentation for details on how to send a message to an Amazon SQS Queue, or review the **describe_nlb_dns.py** in the sample application template package to see how the application template constructs the messages.

Stack Update with VM-Series Auto Scaling Template for AWS (v2.0)

A stack update allows you to modify the resources that the VM-Series Auto Scaling template—firewall-v2.0.template—deploys. Instead of deleting your existing deployment and redeploying the solution, use the stack update to modify the following parameters:

- License—Switch from BYOL to PAYG and vice versa or switch from one PAYG bundle to another.
- Other stack resources— Change the launch configuration parameters such as the Amazon Machine Image (AMI) ID, the AWS instance type, key pair for your auto scaling groups. You can also update the API key associated with the administrative user account on the firewall.

Changing the AMI-ID allows you to deploy new instances of the VM-Series firewalls with a different PAN-OS version.

 ©2024 Palo Alto Networks, Inc.

When you deploy the VM-Series Auto Scaling template, the auto scaling groups and the launch configuration are automatically created for you. The launch configuration is a template that an auto scaling group uses to launch EC2 instance, and it specifies parameters such as the AMI ID, the instance type, key pair for your auto scaling group. To launch VM-Series firewalls with your updated parameters, you must first update the stack and then delete the existing auto scaling groups in each AZ. To prevent service disruption, delete the auto scaling group in one AZ first, and wait for the new firewall instances to launch with the updated stack parameters. Then, verify that the firewalls have inherited the updates you made before you proceed to complete the changes in the other AZ.

For critical applications, perform a stack update during a maintenance window.

You can update stack directly or create change sets. The workflow in this document takes you through the manual stack update.

STEP 1 | In the AWS CloudFormation console, select the parent stack that you want to update and choose **Actions** > **Update Stack**.

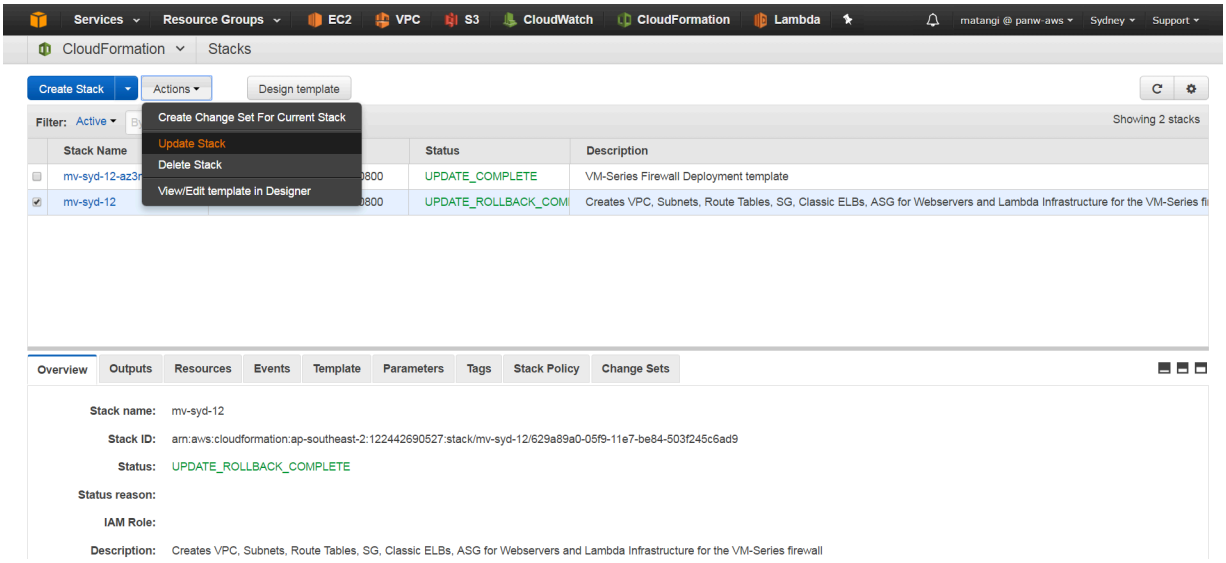

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Modify the resources that you want to update.

- PAN-OS version—To modify the PAN-OS version look up the AMI ID for the version you want to use and enter the ID.
- License option—Switch from BYOL to PAYG or across PAYG bundles 1 and 2.

If you're switching to BYOL, make sure to include the auth code in the bootstrap package (See steps 3 and 5).

If you're switching between PAYG bundle version 1 and 2, look up the AMI ID for the VM-Series firewall.

- Other stack resources— You can modify the AMI ID, the instance type, security group, key pair for the stack resources, or the API key associated with the administrative user account on the firewall.

If you create a new administrative user account or modify the credentials of the existing administrator on the firewall, in order to update that stack and deploy new firewalls with the updated API key, you need to follow the workflow in Modify Administrative Account and Update Stack.

STEP 3 | Acknowledge the notifications and review the changes and click **Update** to initiate the stack update.

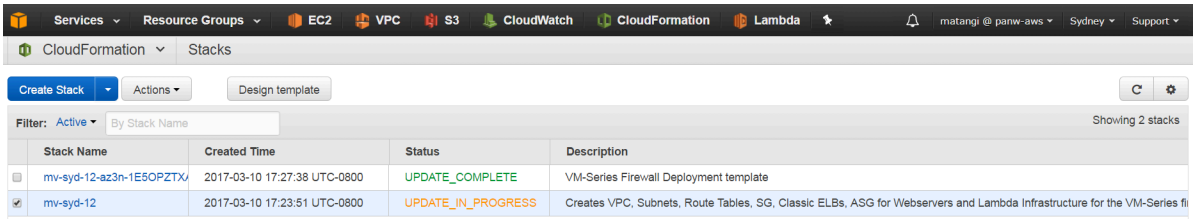

STEP 4 | On the **EC2 dashboard** > **Auto Scaling Groups** and pick an AZ in which to delete the ASG.

Deleting an ASG automatically triggers the process of redeploying a new ASG. The firewalls in the new ASG use the updated stack configuration.

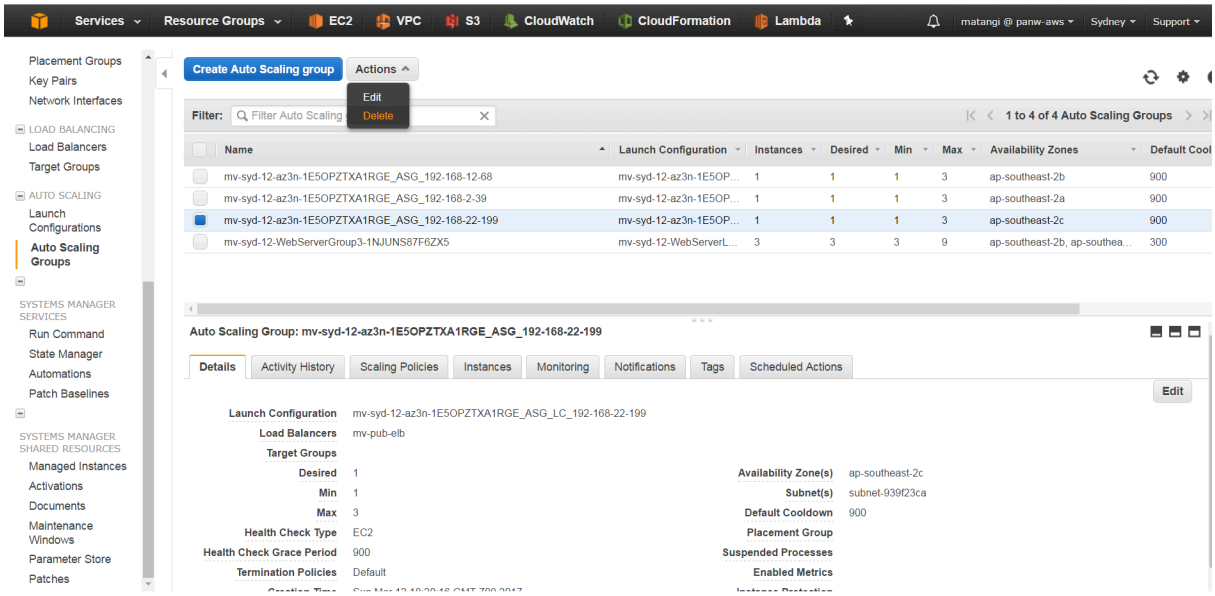

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Verify that the updated parameters are used to launch the VM-Series firewalls in the new ASG.

Use a phased rollout process, where you test the new ASG thoroughly and ensure that the firealls are properly handling traffic. Then, wait one hour before continuing to the next ASG.

STEP 6 | Repeat steps 4 and 5 to replace the ASG in the other AZ.

Modify Administrative Account and Update Stack (v2.0)

If you have already deployed the template and now want to change the password for the administrative account or create a new administrative user account on the VM-Series firewall, you must generate a new API key and update the template stack with the new API key for the administrative user account. And in order to ensure that new firewall instances are configured with the updated administrative user account, you need to export the firewall configuration and rename it to bootstrap.xml, then upload it to the S3 bootstrap folder that the VM-Series AutoScaling template uses.

STEP 1 | Log in to the web interface of the firewall and change the credentials for an existing administrative user or create a new account.

STEP 2 | Generate the API key.

STEP 3 | Export the current running configuration and rename it to bootstrap.xml.

STEP 4 | Upload this bootstrap.xml file to the S3 bootstrap folder; see Customize the Bootstrap.xml File (v2.0).

STEP 5 | Update the API key in the stack to ensure that newly launched firewalls will have the updated administrator account.

See Stack Update with VM-Series Auto Scaling Template for AWS (v2.0).

VM-Series Auto Scaling Templates for AWS Version 2.1

The VM-Series Auto Scaling templates enable you to deploy a single auto scaling group (ASG) of VM-Series firewalls to secure inbound traffic from the internet to your application workloads on AWS. You can deploy the VM-Series firewall ASG and the application workloads within a single VPC as shown below.

 ©2024 Palo Alto Networks, Inc.

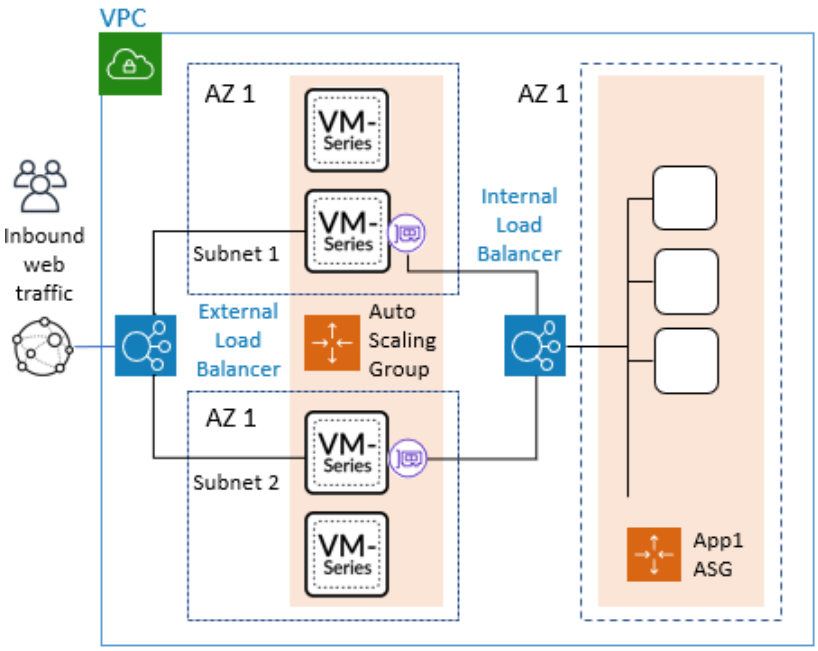

You can also deploy the firewall ASG in a centralized VPC and your application workloads in separate VPCs within the same region, forming a hub and spoke architecture, as shown below.

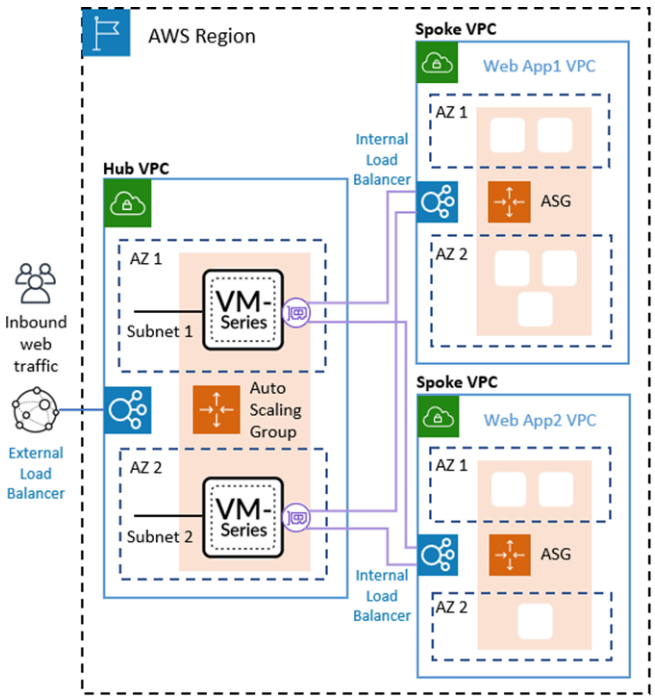

With the hub and spoke architecture you can streamline the delivery of centralized security and connectivity for AWS deployments with many applications, VPCs, or accounts. This architecture can increase agility. Your network security administrators manage the firewall VPC, and DevOps administrators or application developers can manage the application VPCs.

Ensure that the application VPCs connected to the firewall VPC, do not have an Internet Gateway (IGW), and use a continuous monitoring and security compliance service such as Prisma Public Cloud.

 ©2024 Palo Alto Networks, Inc.

You can use a single AWS account or multiple AWS accounts to monitor and secure traffic between VPCs and the internet. Centralizing firewalls in a single VPC can reduce costs for deployments with multiple VPCs and/or multiple accounts.

To provide flexibility with securing your application workloads, version 2.1 allows you to deploy an application load balancer or a network load balancer for both the external load balancer that fronts your VM-Series firewall ASG, and the internal load balancer (ILB) that fronts your application workloads.

When an application load balancer fronts the application workloads, you can connect the firewall VPC to the application VPC using VPC peering. When an NLB fronts the application workloads you can use VPC Peering or an AWS Private Link to connect the firewall and application VPCs, as summarized below:

Firewall VPC LB(External)	Application VPC LB (Internal)	Connection Method
ALB	NLB	AWS Private Link
NLB	NLB	AWS Private Link
NLB	ALB	VPC Peering
ALB	ALB	VPC Peering

If you deploy in a single VPC you can use all the load balancing combinations in the previous table.

You can deploy the templates in both and greenfield (new VPC and applications) and brownfield (existing VPC and applications) use cases.

Template	New	Existing
Firewall	firewall-new-vpc-v2.1.template panw-aws-same-vpc-v2.1.template	firewall-existing-vpc-v2.1.template panw-aws-same-vpc-v2.1.template
Application	panw-aws-nlb-new-vpc-v2.1.template panw-aws-alb-new-vpc-v2.1.template	panw-aws-alb-existing-vpc-v2.1.template panw-aws-nlb-existing-vpc-v2.1.template

What Components Do the VM-Series Auto Scaling Template for AWS (v2.1) Leverage?

The VM-Series Auto Scaling template for AWS includes the following building blocks.

- VM-Series Firewall Templates
- Application Templates
- Lambda Functions
- Panorama

 ©2024 Palo Alto Networks, Inc.

- Bootstrap Files

VM-Series Firewall Templates

The firewall templates deploy an internet-facing external load balancer and VM-Series firewalls within an auto scaling group that spans a minimum of two Availability Zones (AZs). The external load balancer distributes incoming VPC traffic across the pool of VM-Series firewalls. It can be an application load balancer (ALB) or a network load balancer (NLB). The VM-Series firewalls automatically publish custom PAN-OS metrics that enable auto scaling.

Template	Description
firewall-new-vpc-v2.1.template	Deploys a firewall stack with two to four availability zones in a new VPC.
firewall-existing-vpc-v2.1.template	Deploys a firewall stack with two to four availability zones in an existing VPC. To deploy in an existing VPC you must enter: • VPC ID • Internet Gateway ID. This is an existing gateway. • Subnet CIDR lists for the Management, Untrust, Trust, NAT Gateway and Lambda subnets. The template uses the CIDRs to create these subnets. If you choose to create a new ELB, the template connects the firewall ASG to the ELB backend pool. If you use an existing ELB, you must manually connect the firewall ASG to the existing load balancer backend.

See Customize the Firewall Template Before Launch (v2.0 and v2.1) for more on these parameters.

Application Templates

The application template deploys an internal load balancer (ILB) and one auto scaling group with a web server in each availability zone (AZ).

Template	Description
panw-aws-same-vpc-v2.1.template	Deploy application in same VPC as the firewall VPC. You can choose a network or application load balancer.
panw-aws-alb-new-vpc-v2.1.template	Deploy application in a new VPC, using ALB as the internal load balancer, and using VPC Peering between the firewall VPC and application VPC. Supports both same account and cross-account deployments. You must supply the following parameters:

 ©2024 Palo Alto Networks, Inc.

Template	Description
	• Hub account ID • Hub VPC ID for VPC peering • Hub VPC trust subnet CIDRs. The template uses these for route table construction after VPC peering is established, one CIDR per availability zone. • StsAssumeRoleARN (output from the Hub template for SQS access)
panw-aws-nlb-new-vpc-v2.1.template	Deploy application in a new VPC, using NLB as the internal load balancer, and using NLB Endpoint Services/Interfaces to communicate between the firewall VPC and application VPC. You must supply these parameters. • Hub account ID • StsAssumeRoleARN (output from the Hub template for SQS access)
panw-aws-alb-existing-vpc-v2.1.template	Deploy ALB in an existing Application VPC. You must supply the VPC ID for your application, and an existing Subnet ID. This template deploys the load balancer in the application VPC and establishes the lambda resources. You must detach your target workload from any existing load balancer, and connect it to the new load balancer.
panw-aws-nlb-existing-vpc-v2.1.template	Deploy NLB in an existing Application VPC. Deploy application in a new VPC, using NLB as the internal load balancer, and using NLB Endpoint Services/Interfaces to communicate between the firewall VPC and application VPC.

Lambda Functions

AWS Lambda provides robust, event-driven automation without the need for complex orchestration software. AWS Lambda monitors a Simple Queue Service (SQS) to learn about load balancers (ALBs or NLBs) that publish to the queue. When the Lambda function detects a new load balancer, it creates a new NAT policy rule and applies it to the VM-Series firewalls within the ASG. The firewalls have a NAT policy rule for each application, and the firewalls use the NAT policy rule (that maps the port to the load balancer IP address) to forward traffic to the load balancer in front of the application web servers.

The Lambda functions also delete all the configuration items that Lambda added to the device group and template stack in Panorama. This includes the NAT rule, Address Object, and Static

 ©2024 Palo Alto Networks, Inc.

Routes that were pushed to the VM-Series firewall. The Lambda function handles delicensing as well.

To learn more about the Lambda functions, refer to the Palo Alto Networks AWS AutoScale Documentation.

Panorama

You must have Panorama management server in Panorama mode to configure Auto Scaling v2.1.

The Panorama management server provides centralized monitoring and management of multiple Palo Alto Networks next-generation firewalls from a single location. Panorama allows you to oversee all applications, users, and content traversing your network, and use this knowledge to create application enablement policies that protect and control the network. If you are not familiar with Panorama please see the Panorama Administrator's Guide.

Managed firewalls are bootstrapped with an `init-config.txt` file. A sample file is included in the GitHub repository so that you can copy the configuration from the template stack and device group when you create them in your existing Panorama.

The untrust and trust zones created in Panorama must be all lower case.

In Panorama you must configure your network interfaces using DHCP.

- Only eth1/1 should automatically create default route trust and untrust zones.
- The Security Policy zones are named `untrust` and `trust`.

  All zone names must be lower case

- The templates configure an Administrator account named `pandemo` and the password `demopassword`.
- Create a virtual router with the naming convention VR-<*TemplateStackName*>. On the virtual router ECMP tab, enable ECMP.
- To set the DNS server address on Panorama, select **Device** > **Setup** > **Services**. Set the **Primary DNS Server** to 169.254.169.253, the **Secondary DNS Server** to 8.8.8.8, and the **FQDN Refresh Time (sec)** to 60. Panorama requires the AWS DNS server IP address to resolve the FQDN of the internal load balancer on AWS. The FQDN refresh time is the interval at which Panorama commits newly detected internal load balancers.

After the application template has launched, Lambda populates the following in Panorama:

- NAT policy
- Address object for LB in Application Template
- Static routes in the virtual router
- Tcp81 service object

The v2.1 firewall template includes an AWS NAT gateway that the firewalls use to initiate outbound requests for retrieving updates, connecting to Panorama, and publishing metrics to AWS CloudWatch. The NAT Gateways also have Elastic IP addresses attached to them for each zone.

 ©2024 Palo Alto Networks, Inc.

You need the following Panorama resources to work with the Auto Scale templates for AWS.

Panorama API Key	You need a Panorama API key to authenticate the API. Lambda uses your API key to autoconfigure template and device group options. To generate the API key, see Get Your API Key.
Panorama License Deactivation Key	The template requires a license deactivation API key and the "Verify Update Server Identity" to be enabled to deactivate the license keys from Panorama. The license deactivation key should be obtained from Palo Alto Customer Support Portal as described in Install a License API Key.
Panorama VM-Auth-Key	You need a vm-auth-key to enable bootstrapped firewalls to connect to Panorama and receive their bootstrap configuration. See Generate the VM Auth Key on Panorama.
Panorama Management Interface Access	• **Port 443 (HTTPS)**—Upon initial deployment of the firewall template, leave HTTPS open so Lambda can connect to Panorama. Wait to receive the following confirmation of connection in Panorama: 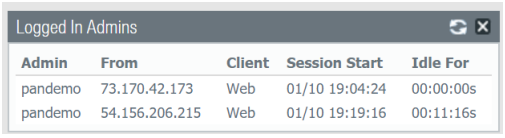 Logged In Admins Admin From Client Session Start Idle For pandemo 73.170.42.173 Web 01/10 19:04:24 00:00:00s pandemo 54.156.206.215 Web 01/10 19:19:16 00:11:16s When you secure port 443 you specify an IP address range from which you will allow connections, as well as the EIPs assigned to the NAT gateways. There are two NAT gateways and the EIPs associated with them. To find NAT gateway EIPs in AWS, go to **VPC** > **NAT Gateways**. Note the EIP information for the security group for HTTPS. • **Port 3978**—Port 3978 must be able to receive traffic from any IP address.

Bootstrap Files

The GitHub auto scaling repository includes an `init-cfg.txt` file so that the VM-Series firewall has the basic configuration to:

- Perform interface swap so the VM-Series firewall `untrust` traffic uses AWS ENI for eth0.
- Communicate to Panorama for device group and template configuration.

The auto scaling GitHub repository has the basic configuration to get started. This auto scaling solution requires swapping the dataplane and management interfaces to enable the load balancer to forward web traffic to the VM-Series firewall auto scaling tier. For details on management interface mapping with the Amazon ELB as shown in Management Interface Mapping for Use with Amazon ELB.

 ©2024 Palo Alto Networks, Inc.

Plan to Deploy VM-Series Auto Scaling Templates for AWS (v2.1)

Before starting the deployment, review the following resources.

- See Auto Scaling VM-Series Firewalls with the Amazon ELB Service for an overview of template features, and account planning.
- Customize the Firewall Template Before Launch (v2.0 and v2.1). The basic parameters in this topic apply to all template versions.
- How Does the VM-Series Auto Scaling Template for AWS (v2.0 and v2.1) Enable Dynamic Scaling?

  These concepts apply to all template versions.

Launch the Firewall Template (v2.1)

You can choose to deploy the firewall and application templates in the same VPC or in separate VPCs.

The templates support a hub and spoke architecture in which you can deploy the firewall template in one AWS account and use it as a hub to secure applications (spokes) that belong to the same or to different AWS accounts.

This workflow tells you how to deploy the external load balancer and the VM-Series firewalls using the firewall template. The vm-auth-key must be configured on Panorama prior to launching this template.

STEP 1 | Review the checklists in Plan to Deploy VM-Series Auto Scaling Templates for AWS (v2.1) and Plan the VM-Series Auto Scaling Template for AWS (v2.0 and v2.1).

Verify that you have completed the following tasks:

- (For PAYG only) Review and accept the EULA for the PAYG bundle you plan to use.
- (For BYOL only) Obtain the auth code for a bundle that supports the number of firewalls that might be required for your deployment. You must save this auth code in a text file named `authcodes` (no extensions), and put the `authcodes` file in the `/license` folder of the bootstrap package.

  If you use individual auth codes instead of a bundle, the firewall only retrieves the license key for the first auth code in the file.
- Download the files required to launch the VM-Series Auto Scaling v2.1 template from the GitHub repository.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Modify the `init-cfg.txt` file and upload it to the `/config` folder.

Because you use Panorama to bootstrap the VM-Series firewalls, your `init-cfg.txt` file should be modified as follows. No `bootstrap.xml` file is needed.

```
type=dhcp-client
```

```
ip-address=
```

```
default-gateway=
```

```
netmask=
```

```
ipv6-address=
```

```
ipv6-default-gateway=
```

```
hostname=
```

```
vm-auth-key=
```

```
panorama-server=
```

```
panorama-server-2=
```

```
tplname=AWS-tmplspoke1
```

```
dgname=AWS-dgspoke1
```

```
dns-primary=169.254.169.253
```

```
dns-secondary=8.8.8.8
```

```
op-command-modes=mgmt-interface-swap
```

```
dhcp-send-hostname=yes
```

 ©2024 Palo Alto Networks, Inc.

```
dhcp-send-client-id=yes
```

```
dhcp-accept-server-hostname=yesdhcp-accept-server-domain=yes
```

```
vm-series-auto-registration-id=
```

```
vm-series-auto-registration-pin-value=
```

Verify that **`op-command-modes=mgmt-interface-swap`** exists. This is the command for swapping the management interface (mgmt) and the dataplane interface (ethernet 1/1) on the VM-Series firewall on AWS. Use the AWS DNS server IP address of 169.254.169.253 for faster load balancer DNS name resolution.

You must add the device certificate auto-registration PIN to the init-cfg.txt file to automatically install a device certificate when your VM-Series firewall instance is deployed.

STEP 3 | (For BYOL only) Add the license auth code in the `/license` folder of the bootstrap package.

1. Use a text editor to create a new text file named **`authcodes`** (no extension).
2. Add the authcode for your BYOL licenses to this file, and save. The authcode must represent a bundle, and it must support the number of firewalls that might be required for your deployment. If you use individual authcodes instead of a bundle, the firewall only retrieves the license key for the first authcode in the file.

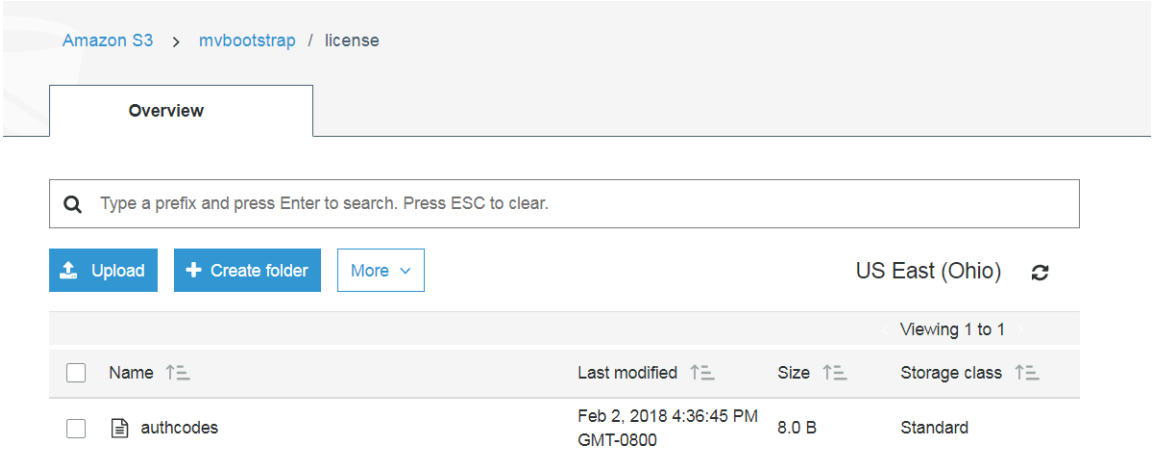

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Upload Lambda code for the firewall template (`panw-aws-zip`) and the Application template (`ilb.zip`) to an S3 bucket. You can use the same S3 bucket that you use for bootstrapping.

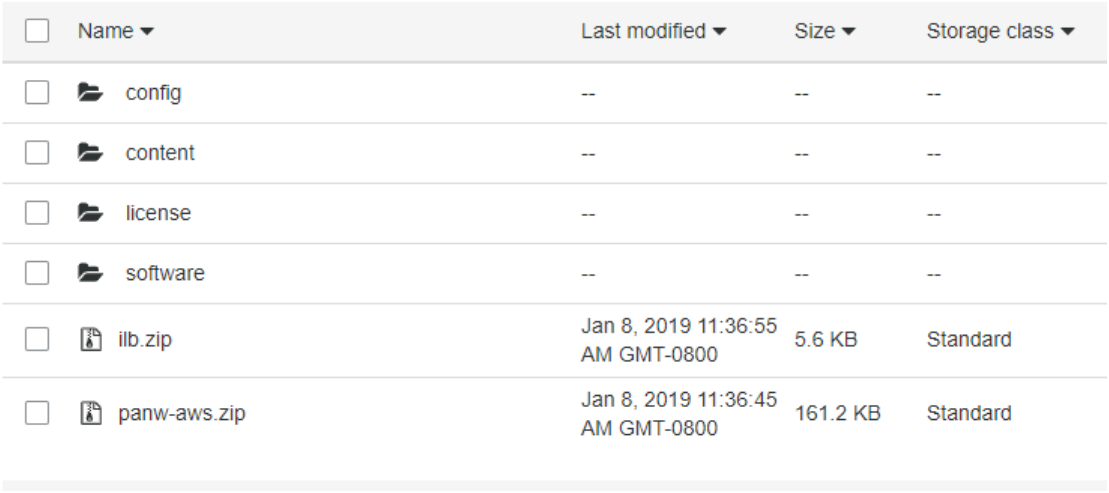

Name	Last modified	Size	Storage class
config	--	--	--
content	--	--	--
license	--	--	--
software	--	--	--
ilb.zip	Jan 8, 2019 11:36:55 AM GMT-0800	5.6 KB	Standard
panw-aws.zip	Jan 8, 2019 11:36:45 AM GMT-0800	161.2 KB	Standard

If the Application stack is managed by a different account than the firewall, use the Application account to create another s3 bucket in the same AWS region as the firewall template and copy `ilb.zip` to that s3 bucket.

STEP 5 | Select the firewall template.

1. In the AWS Management Console, select **CloudFormation** > **Create Stack**.
2. Select Upload a template to Amazon S3, to choose the application template to deploy the resources that the template launches within the same VPC as the firewalls, or to a different VPC. Click Open and Next.
3. Specify the Stack name. The stack name allows you to uniquely identify all the resources that are deployed using this template.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Configure the parameters for the VPC.

1. Be sure to select at least two availability zones

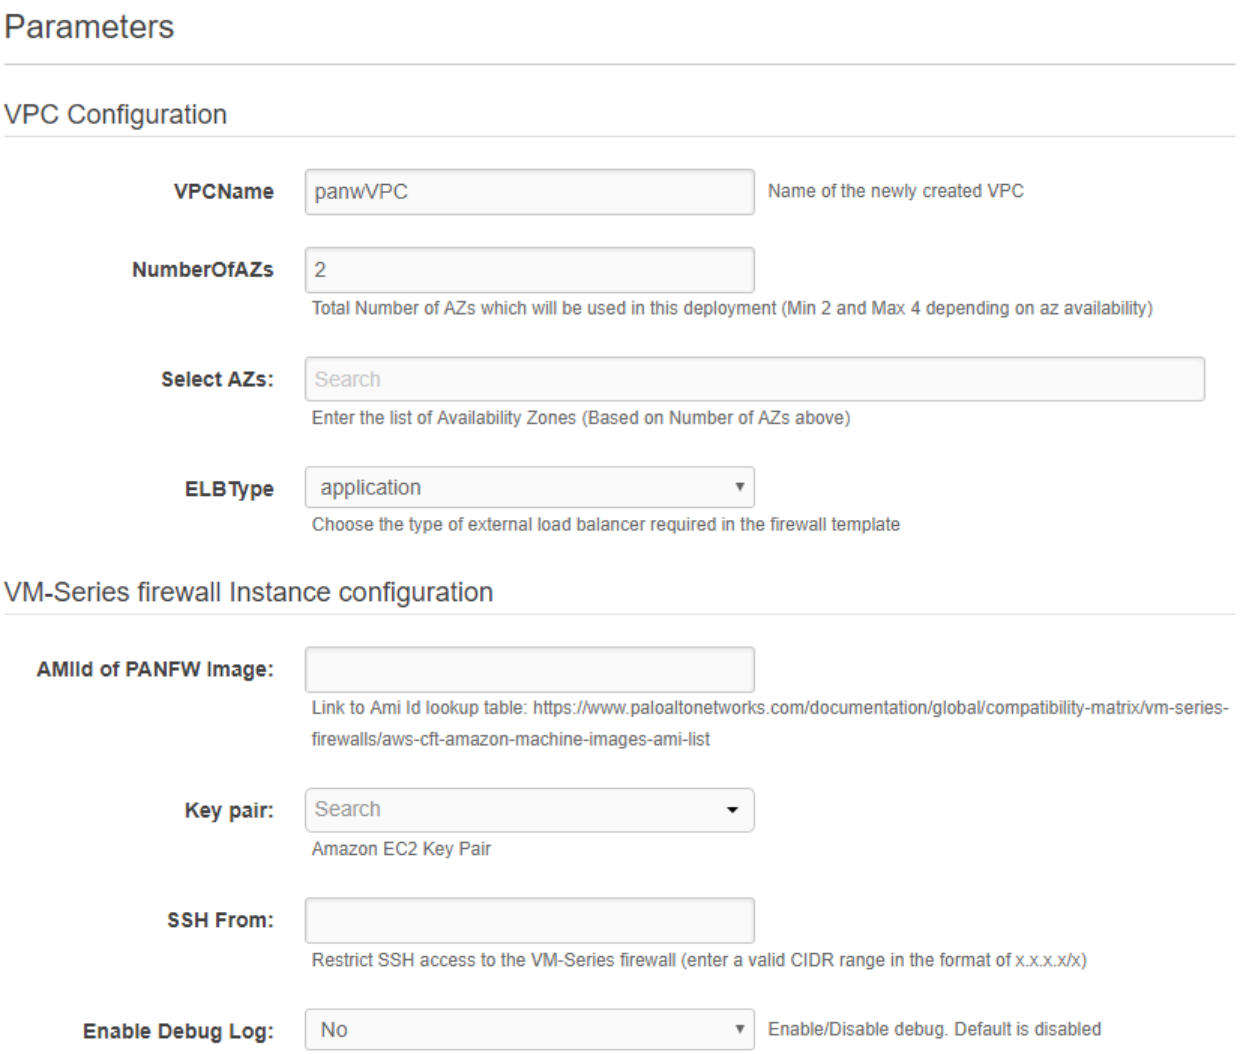

2. Look up the AMI ID for the VM-Series firewall and enter it. Make sure that the AMI ID matches the AWS region, PAN-OS version and the BYOL or PAYG licensing option you opted to use.
3. Select the EC2 **Key pair** (from the drop-down) for launching the firewall. To log in to the firewalls, you must provide the name of this key pair and the private key associated with it.
4. For the **SSH from field**, the firewalls will be managed by Panorama and do NOT have an EIP for the management interface. But just in case you decide to assign an EIP configure the IP range you would connect from.
5. Select **Yes** if you want to **Enable Debug Log**. Enabling the debug log generates more verbose logs that help with troubleshooting issues with the deployment. These logs are generated using the stack name and are saved in AWS CloudWatch.

By default, the template uses CPU utilization as the scaling parameter for the VM-Series firewalls. Custom PAN-OS metrics are automatically published to the CloudWatch namespace that matches the stack name you specified earlier.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Specify the name of the Amazon S3 bucket(s).

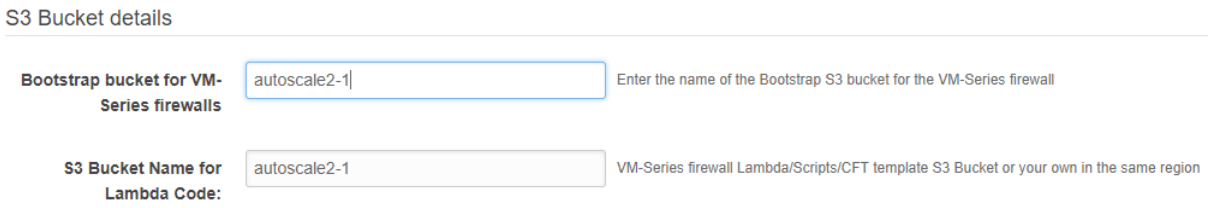

1. Enter the name of the S3 bucket that contains the bootstrap package.

   If the bootstrap bucket is not set up properly or if you enter the bucket name incorrectly, the bootstrap process fails, and you cannot log in to the firewall. Health checks for the load balancers also fail.
2. Enter the name of the S3 bucket that contains the panw-aws.zip file. As mentioned earlier you can use one S3 bucket for the Bootstrap and Lambda code.

STEP 8 | Specify the keys for enabling API access to the firewall and Panorama.

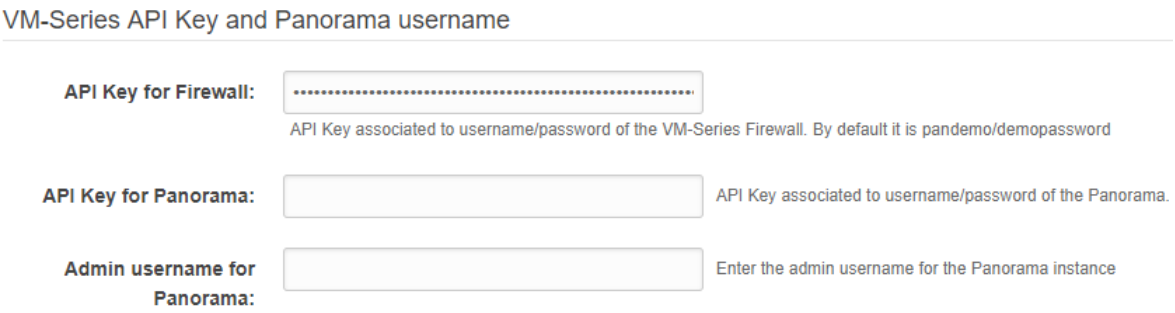

1. Enter the key that the firewall must use to authenticate API calls. The default key is based on the sample file and you should only use it for testing and evaluation. For a production deployment, you must create a separate PAN-OS login just for the API call and generate an associated key.
2. Enter the API Key to allow AWS Lambda to make API calls to Panorama. For a production deployment, you should create a separate login just for the API call and generate an associated key.

STEP 9 | Enter the name for the application load balancer.

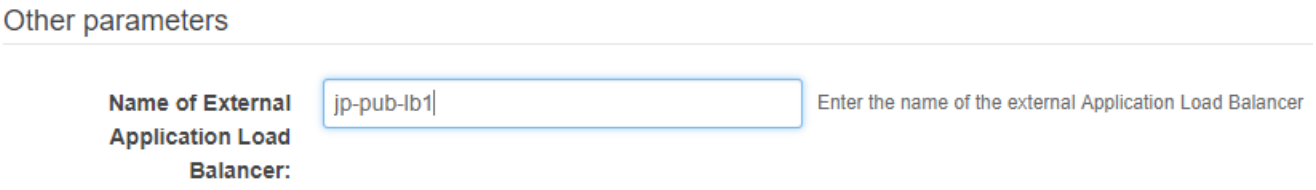

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Review the template settings and launch the template.

1. Select **I acknowledge that this template might cause AWS CloudFormation to create IAM resources.**
2. Click **Create** to launch the template. The CREATE_IN_PROGRESS event displays.
3. On successful deployment the status updates to CREATE_COMPLETE.

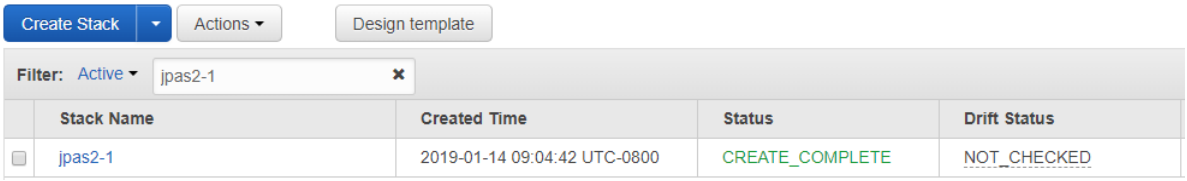

STEP 11 | Verify that the template has launched all required resources.

1. On the EC2 Dashboard, select Auto Scaling Groups. Verify that in each AZ, you have one ASG for the VM-Series firewalls. The ASG name prefix includes the stack name.
2. On the AWS Management Console, select the stack name to view the Output for the list of resources.
3. Your output should look similar to the output in the following image.
   - Take note of the Network Load Balancer Queue name.
   - Take note of the Elastic Load Balancer public DNS name.

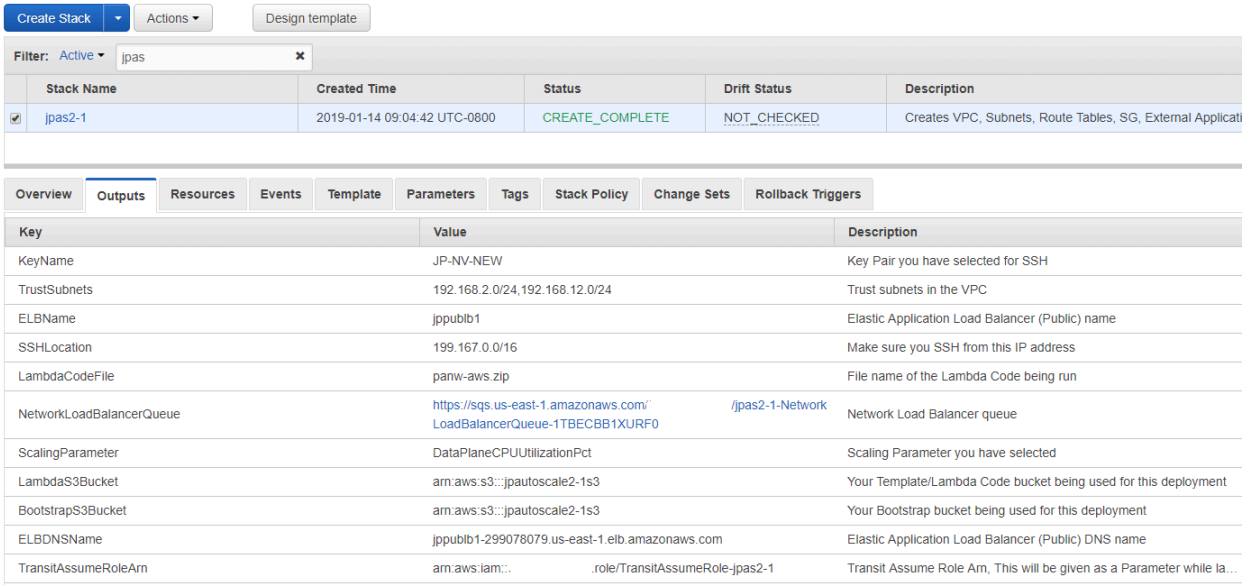

It may take up to 20 minutes for the firewalls to boot up and be available to handle traffic.

When you are finished with a testing or a production deployment, the only way to ensure charges stop occurring is to completely delete the stack. Shutting down instances, or changing the ASG maximum to 0 is not sufficient.

 ©2024 Palo Alto Networks, Inc.

STEP 12 | Save the following firewall template information. You must provide these values as inputs when deploying the application template.

- **IP addresses of the NAT Gateway in each AZ**—You need this IP address to restrict HTTPS access to your Panorama so that Lambda can use the EIPs for the NAT Gateway to communicate with Panorama when needed.
- **Network Load Balancer SQS URL**—A Lambda function in the firewall stack monitors this queue so that it can learn about any network load balancers that you deploy and create NAT policy rules (one per application) in the Panorama that enable the firewalls to send traffic to the network load balancer IP address.

Launch the Application Template (v2.1)

The application templates allow you to complete the sandwich topology and are provided so that you can evaluate the auto scaling solution. This application template deploys either an application or network load balancer and a pair of web servers behind the auto scaling group of VM-Series firewalls, which you deployed using the firewall template.

Use this template to evaluate the solution but customize your own template to deploy to production. For a custom template, make sure to enable SQS messaging between the Application template and the Firewall template.

When launching the application template, you must select the template based on whether you want to deploy the application template within the same VPC in which you deployed the firewall template or in a separate VPC. See Enable Traffic to the ELB Service (v2.0 and v2.1).

STEP 1 | Create an S3 bucket from which you will launch the application template.

- If this is a cross-account deployment, create a new bucket.
- If there is one account you can create a new bucket or use the S3 bucket you created earlier (you can use one bucket for everything).

STEP 2 | Upload the ilb.zip file into the S3 bucket.

Name	Last modified	Size	Storage class
config	--	--	--
content	--	--	--
license	--	--	--
software	--	--	--
ilb.zip	Jan 8, 2019 11:36:55 AM GMT-0800	5.6 KB	Standard
panw-aws.zip	Jan 8, 2019 11:36:45 AM GMT-0800	161.2 KB	Standard

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Select the application launch template you want you launch.

1. In the AWS Management Console, select **CloudFormation** > **CreateStack**
2. Select Upload a template to Amazon S3, to choose the application template to deploy the resources that the template launches within the same VPC as the firewalls, or to a different VPC. Click **Open** and **Next**.
3. Specify the Stack name. The stack name allows you to uniquely identify all the resources that are deployed using this template.

STEP 4 | Configure the parameters for the VPC and network load balancer.

1. Select the two Availability Zones that your setup will span in Select list of AZ. If you are deploying within the same VPC make sure to select the same Availability Zones that you selected for the firewall template.
2. If deploying to a new VPC enter a CIDR Block for the VPC. The default CIDR is 192.168.0.0/16.
3. If deploying to the same VPC you will select the previous VPC and use the Trust subnets.

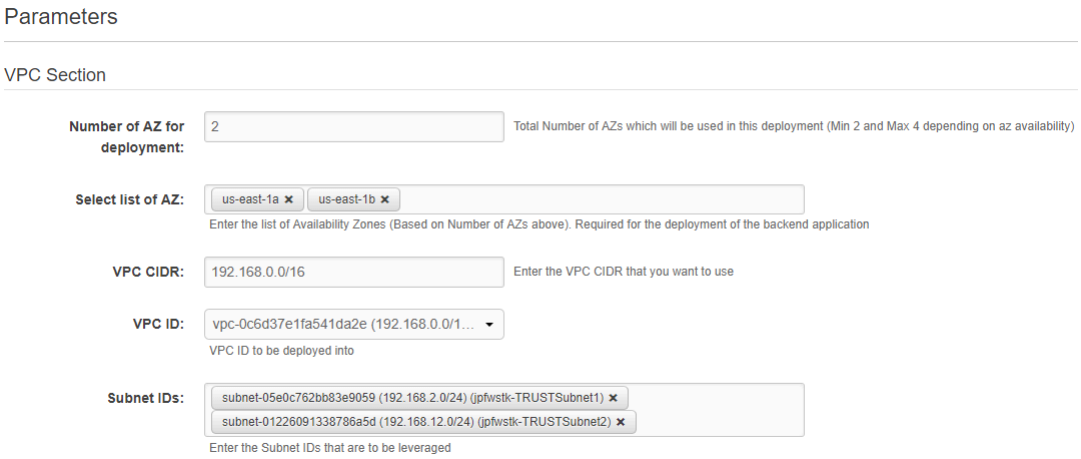

STEP 5 | Select the load balancer type.

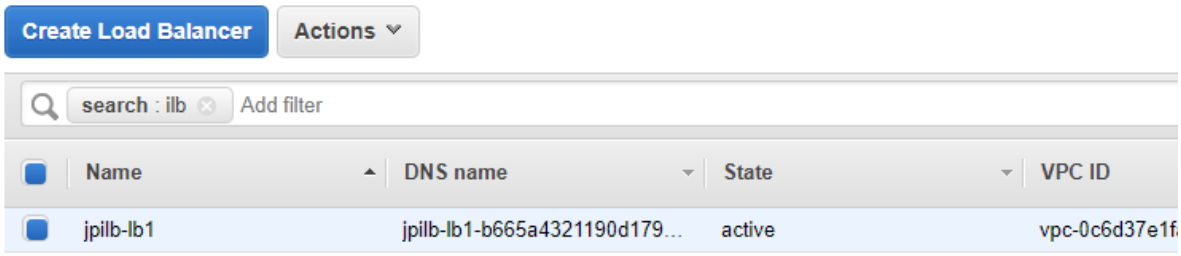

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Configure the parameters for Lambda.

1. Enter the S3 bucket name where ilb.zip is stored.
2. Enter the name of the zip file name.
3. Paste the SQS URL that you copied earlier.

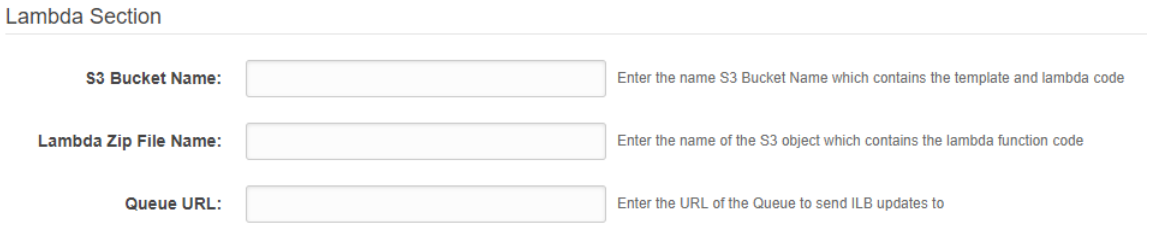

STEP 7 | Modify the web server EC2 instance type to meet your needs.

STEP 8 | Select the EC2 Key pair (from the drop-down) for launching the web servers. To log in to the web servers, you must provide the key pair name and the private key associated with it.

STEP 9 | Select the IP address of the network you will be accessing the servers from for management access only. Web traffic comes through the ELBDNS name you copied when you launched the firewall template.

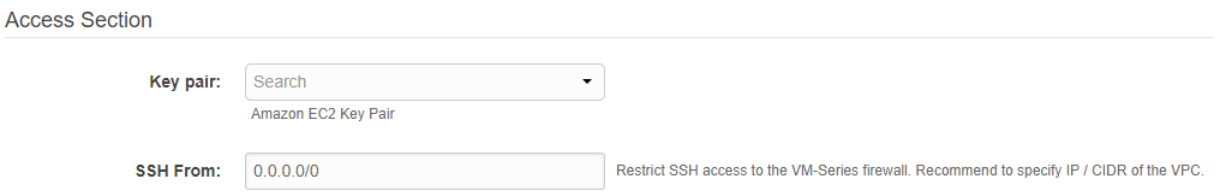

STEP 10 | Review the template settings and launch the template.

 ©2024 Palo Alto Networks, Inc.

STEP 11 | After completion of the application template it can take up to 20 minutes for the web pages to become active.

1. Verify that the application template load balancer is marked active.

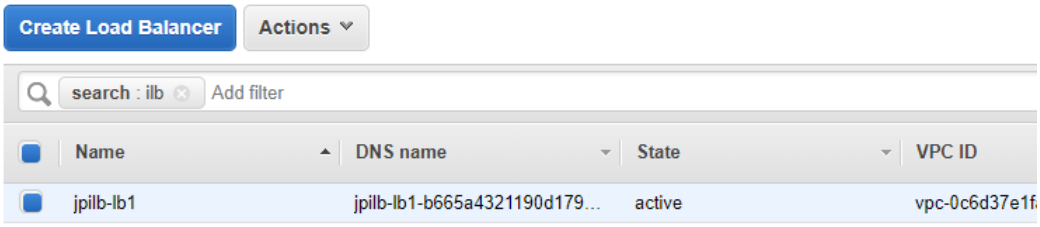

2. Verify that Panorama has a NAT object in the device group.

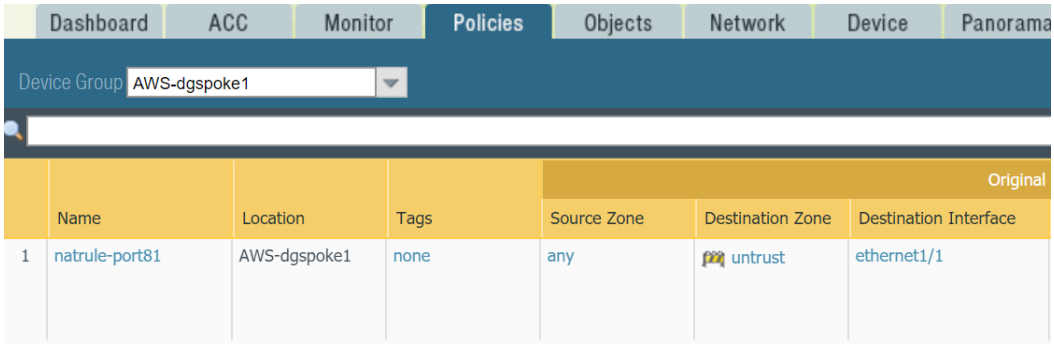

3. Verify that Panorama has an address object in the device group.

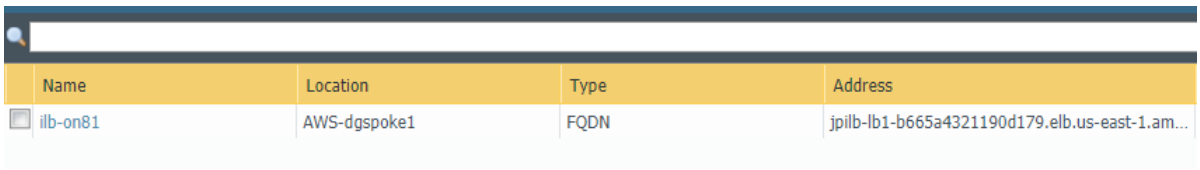

4. Verify that Panorama has static routes in the template stack.

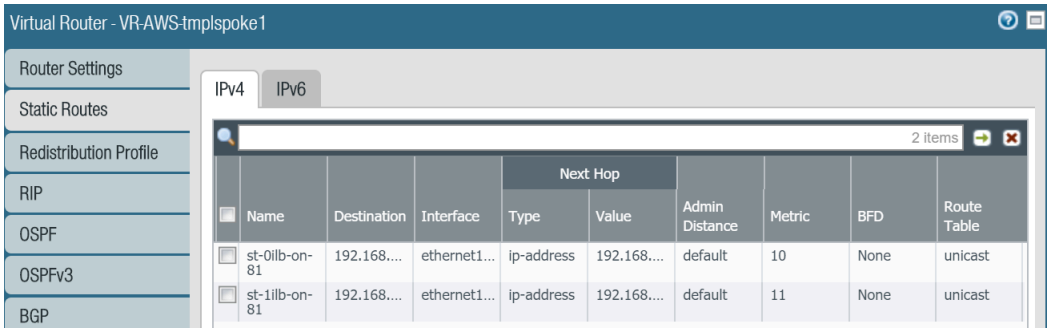

STEP 12 | Get the DNS name you saved earlier for the application load balancer and enter it into a web browser.

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Upon successful launch your browser should look like this output.

Congratulations, you have successfully launched VM-Series ASG CloudFormation. This file is coming from Webserver Region: us-east-1

StackID: arn:aws:cloudformation:us-east-1: :stack/jpappstk1/1c197cb0-18f4-11e9-b92e-0ab87eb901cc

StackName: jpappstk1

Create a Custom Amazon Machine Image (v2.1)

A custom VM-Series AMI gives you the consistency and flexibility to deploy a VM-Series firewall with the PAN-OS version you want to use on your network instead of being restricted to using only an AMI that is published to the AWS public Marketplace or to the AWS GovCloud Marketplace. Using a custom AMI speeds up the process of deploying a firewall with the PAN-OS version of your choice because it reduces the time to provision the firewall with an AMI published on the AWS public or AWS GovCloud marketplace, and then perform software upgrades to get to the PAN-OS version you want to use on your network. Additionally, you can use the custom AMI in the Auto Scaling VM-Series Firewalls CloudFormation Templates or any other templates that you have created.

You can create a custom AMI with the BYOL, Bundle 1, or Bundle 2 licenses. The process of creating a custom AMI requires you to remove all configuration from the firewall and perform a private data reset, so in this workflow you'll launch a new instance of the firewall from the AWS Marketplace instead of using an existing firewall that you have fully configured.

When creating a custom AMI with a BYOL version of the firewall, you must first activate the license on the firewall so that you can access and download PAN-OS content and software updates to upgrade your firewall, and then deactivate the license on the firewall before performing the private data reset and creating the custom AMI. If you do not deactivate the license, you lose the license that you applied on this firewall instance.

STEP 1 | Launch the VM-Series firewall from the Marketplace.

See Launch the VM-Series firewall.

STEP 2 | (Only for BYOL) Activate the license.

STEP 3 | Install latest content on the firewall.

STEP 4 | (Only for BYOL) Deactivate the license.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Perform a private data reset.

A private data reset removes all logs and restores the default configuration.

The system disks are not erased, so the content updates from Step 4 are intact.

1. Access the firewall CLI.
2. Export a copy of the configuration.
3. Remove all logs and restore the default configuration.

   ```
   request system private-data-reset
   ```

   Enter **y** to confirm.

   The firewall reboots to initialize the default configuration.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Create the custom AMI.

1. Log in to the AWS Console and select the EC2 Dashboard.
2. **Stop** the VM-Series firewall.
3. Select the VM-Series firewall instance, and click **Image** > **Create Image**.

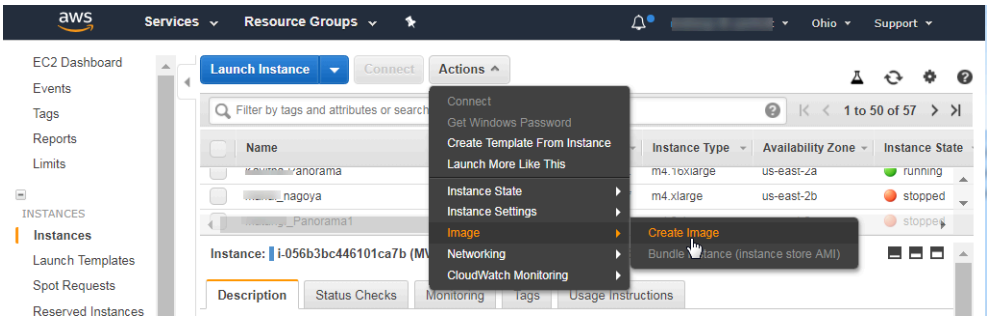

4. Enter a custom image name, and click **Create Image**.

   The disk space of 60GB is the minimum requirement.

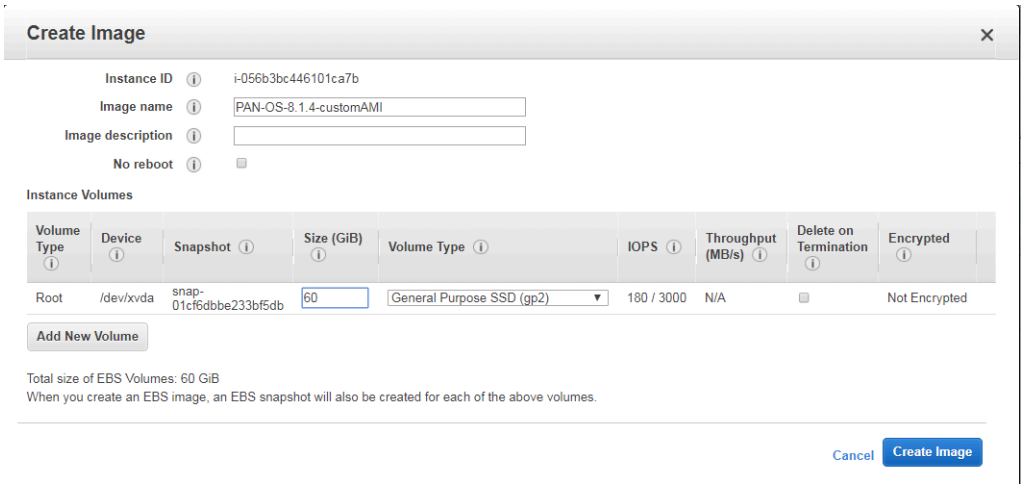

5. Verify that the custom AMI is created and has the correct product code.
   1. On the EC2 Dashboard, select **AMI**.
   2. Select the AMI that you just created. Depending on whether you selected an AMI with the BYOL, Bundle 1, or Bundle 2 licensing options, you should see one of the following **Product Codes** in the details:
      - BYOL—6njl1pau431dv1qxipg63mvah
      - Bundle 1—6kxdw3bbmdeda3o6i1ggqt4km
      - Bundle 2—806j2of0qy5osgjjixq9gqc6g

 ©2024 Palo Alto Networks, Inc.

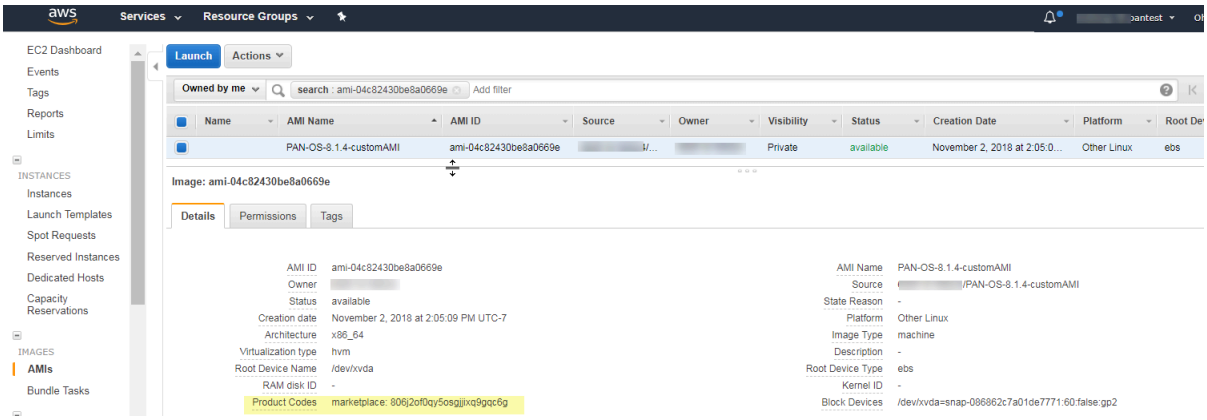

STEP 7 | Encrypt EBS Volume for the VM-Series Firewall on AWS.

If you plan to use the custom AMI with EBS encryption for an Auto Scaling VM-Series Firewalls with the Amazon ELB Service deployment, you must use the default master key for your AWS account.

VM-Series Auto Scaling Template Cleanup (v2.1)

If you deployed the templates as a test, delete them to save resources and lower costs.

STEP 1 | In the AWS Management Console, select **Cloud Formation** > **Create Stack**.

STEP 2 | Locate the firewall template and application template you launched previously and delete both templates.

For more information on deleting template stacks see, "What is AWS CloudFormation?"

Failure to delete your template stack incurs charges from AWS.

SQS Messaging Between the Application Template and Firewall Template (v2.1)

VM-Series firewalls deployed using one of the firewall templates can detect and send traffic to the load balancers to which you want to automatically distribute incoming traffic. To accomplish this, the firewall template includes a lambda function that monitors a Simple Queue Service for messages. The message allows the lambda function to learn about a new load balancer and then automatically create a NAT policy rule on the firewall to send traffic to the load balancer's IP address. To route traffic properly within the AWS infrastructure, the message must also include basic information on the DNS, VPC ID, and the AZ to which the load balancer belongs.

If you are building your own application template, you must set up your application template to post ADD and DEL messages to the SQS URL that the firewall template uses to learn about load balancers to which it must distribute traffic in your environment:

- ADD-NLB message that informs the firewalls when a new network load balancer is available.
- DEL-NLB message that informs the firewalls when a network load balancer has been terminated and is no longer available.
- ADD-ALB message that informs the firewalls when a new application load balancer is available.

 ©2024 Palo Alto Networks, Inc.

- DEL-ALB message that informs the firewalls when a application load balancer has been terminated and is no longer available.

The following examples of each message type include sample values. You must modify these messages with values that match your deployment.

ADD-NLB Message

```
msg_add_nlb= {

       "MSG-TYPE": "ADD-NLB",

       "AVAIL-ZONES": [

       {

           "NLB-IP":"192.168.2.101",

           "ZONE-NAME":"us-east-2a",

           "SUBNET-ID": "subnet-2a566243"

       },

       {

           "NLB-IP":"192.168.12.101",

           "ZONE-NAME":"us-east-2b",

           "SUBNET-ID": "subnet-2a566243 "

       }

       ],

            "DNS-NAME": "publicelb1-2119989486.us-
east-2.elb.amazonaws.com",
```

 ©2024 Palo Alto Networks, Inc.

```
            "VPC-ID": "vpc-42ba9f2b",

            "NLB-NAME": "publicelb1"

}
```

DEL-NLB Message

```
msg_del_nlb= {

 "MSG-TYPE": "DEL-NLB",

 "DNS-NAME": "publicelb1-2119989486.us-east-2.elb.amazonaws.com",

}
```

ADD-ALB

```
{   "AVAIL-ZONES": [

       {

            "SUBNET-CIDR": "172.32.0.0/24",

            "SUBNET-ID": "subnet-0953a3a8e2a8208a9",

            "ZONE-NAME": "us-east-2a"

       },

       {

            "SUBNET-CIDR": "172.32.2.0/24",

            "SUBNET-ID": "subnet-0a9602e4fb0d88baa",

            "ZONE-NAME": "us-east-2c"
```

 ©2024 Palo Alto Networks, Inc.

```
        },
        {
            "SUBNET-CIDR": "172.32.1.0/24",
            "SUBNET-ID": "subnet-0b31ed16f308b3c4d",
            "ZONE-NAME": "us-east-2b"
        }
    ],
    "VPC-PEERCONN-ID": "pcx-0538bb05dbe2e1b8e",
    "VPC-CIDR": "172.32.0.0/16",
    "ALB-NAME": "appILB-908-0",
    "ALB-ARN":"arn:aws:elasticloadbalancing:us-
east-2:018147215560:loadbalancer/app/appILB-908-0/1997ed20eeb5bcef",
    "VPC-ID": "vpc-0d9234597da6d9147",
    "MSG-TYPE": "ADD-ALB",
    "DNS-NAME": "internal-appILB-908-0-484644265.us-
east-2.elb.amazonaws.com"
}
```

DEL-ALB Message

```
{
    "MSG-TYPE": "DEL-ALB",
```

 ©2024 Palo Alto Networks, Inc.

```
   "DNS-NAME": "internal-appILB-908-0-484644265.us-
east-2.elb.amazonaws.com"
```

```
}
```

Refer to the AWS documentation for details on how to send a message to an Amazon SQS Queue.

Stack Update with VM-Series Auto Scaling Template for AWS (v2.1)

A stack update allows you to modify the resources that the VM-Series Auto Scaling template firewall template deploys. Instead of deleting your existing deployment and redeploying the solution, use the stack update to modify launch configuration parameters.

You can modify the AWS instance type, the key pair for your auto scaling groups, and the APi key associated with the adminstrative user account on the firewall.

You do not have to update the stack to modify default notifications or create auto scaling alarms. See Change Scaling Parameters and CloudWatch Metrics (v2.1).

When you deploy the VM-Series Auto Scaling template, the auto scaling groups and the launch configuration are automatically created for you. The launch configuration is a template that an auto scaling group uses to launch EC2 instance, and it specifies parameters such as the instance type, the key pair for your auto scaling group, or the API key associated with the administrative user account on the firewall.

For critical applications, perform a stack update during a maintenance window.

You can update your stack directly or create change sets. The workflow in this document takes you through the manual stack update.

STEP 1 | In the AWS CloudFormation console, select the parent stack that you want to update and choose **Actions** > **Update Stack**.

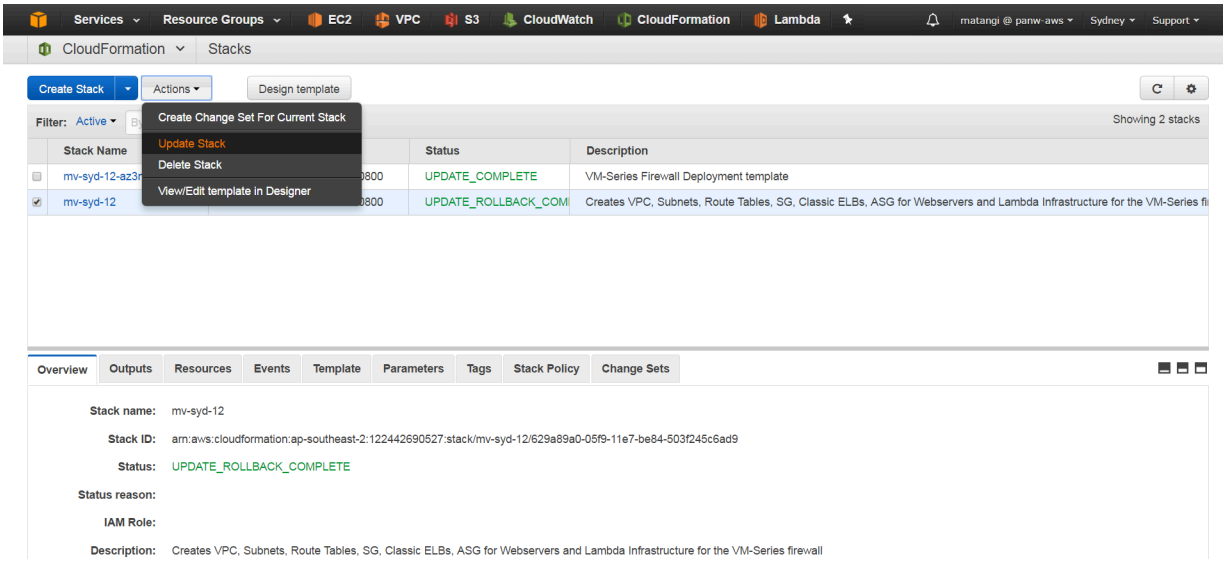

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Modify the resources that you want to update.

You can modify the instance type, security group, key pair for the stack resources, or the API key associated with the administrative user account on the firewall.

If you create a new administrative user account or modify the credentials of the existing administrator on the firewall, in order to update that stack and deploy new firewalls with the updated API key, you need to follow the workflow in Modify Administrative Account (v2.1).

STEP 3 | Acknowledge the notifications and review the changes and click **Update** to initiate the stack update.

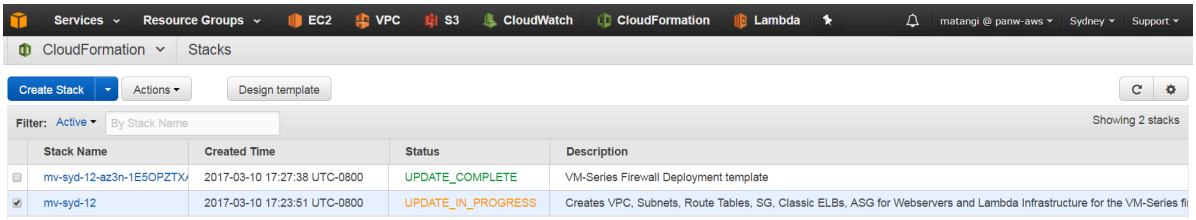

Modify Administrative Account (v2.1)

If you have already deployed the template and now want to change the password for the administrative account or create a new administrative user account on the VM-Series firewall, you must generate a new API key and update the template stack with the new API key for the administrative user account.

STEP 1 | Log in to the web interface of the firewall and change the credentials for an existing administrative user or create a new account.

STEP 2 | Generate the API key.

STEP 3 | Update the API key in the stack to ensure that newly launched firewalls have the updated administrator account.

See Stack Update with VM-Series Auto Scaling Template for AWS (v2.0).

Change Scaling Parameters and CloudWatch Metrics (v2.1)

This task describes how to use custom PAN-OS metrics as scaling parameters to trigger auto scale actions.

When you launch the firewall template, the template creates a namespace with scale-in and scale-out policies that you can use to define auto scaling actions. The policy names include the namespace, as shown below:

- <Custom Namespace>-scalein - Remove 1 instance
- <Custom Namespace>-scaleout - Add 1 instance

Each PAN-OS metric has a default notification that you can delete and replace with auto scale actions. For each metric, create two actions: one that determines when to add a VM-Series firewall, and another that determines when to remove a VM-Series firewall.

STEP 1 | In AWS, select **Services** > **CloudWatch** > **Metrics**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Choose a **Custom Namespace** link, and select the metrics link to view the custom PAN-OS metrics.

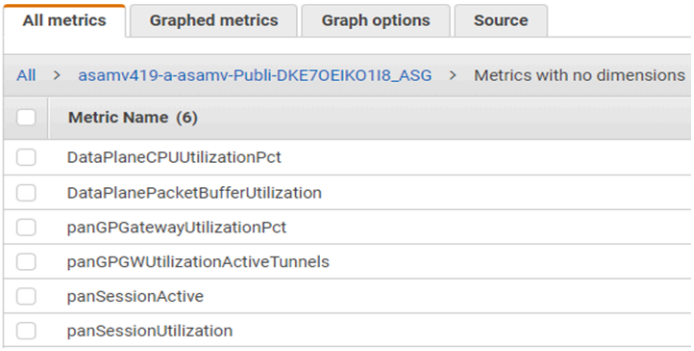

STEP 3 | Check a box to select a metric, then select the **Graphed metrics** tab.

1. In the **Statistics** column, choose a statistic criteria (such as average, minimum and maximum) and choose a time period.
2. In the **Actions** column select the bell (Create alarm).

STEP 4 | Define an alarm that removes a firewall when CPU utilization meets or falls below the criteria you set, over the time frame you set.

1. Select **Edit** to change the graph title.
2. Under **Alarm details** fill in the **Name** and **Description**, choose an operator, and set the minimum value to maintain the current instances. If the minimum value is not maintained, an instance is removed.
3. Under **Actions**, **delete** the default notification.
4. Select **+AutoScaling Action**.
   - Use the **From the** list to select your namespace.
   - From **Take this action**, select the policy to remove an instance.

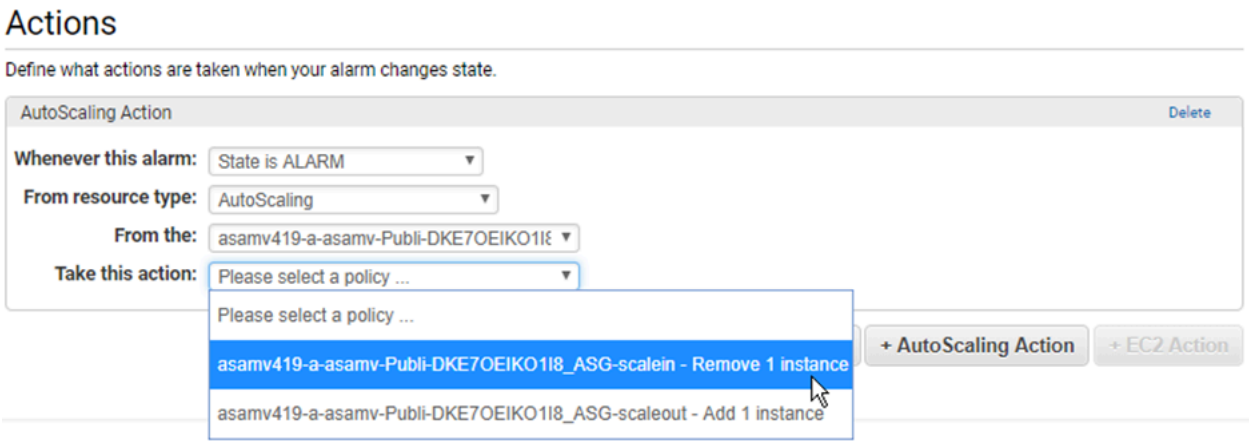

5. Select **Create Alarm**.

STEP 5 | Create a second alarm that adds a firewall when CPU utilization meets or exceeds the criteria you set.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | To view your alarms, select **Services** > **CloudWatch** > **Alarms**.

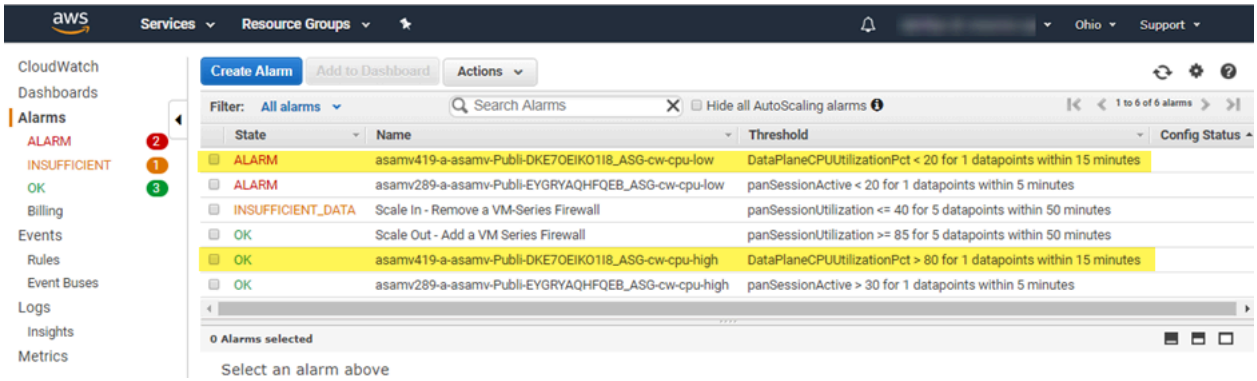

To edit an alarm from this window, check the box next to the alarm and select **Action** > **Edit**.

 ©2024 Palo Alto Networks, Inc.

List of Attributes Monitored on the AWS VPC

As you provision or modify virtual machines in your AWS VPCs, you have two ways of monitoring these instances and retrieving the tags for use as match criteria in dynamic address groups.

- **AWS Resource Information**—On a next-gen firewall, you can monitor up to a total of 32 tags—14 pre-defined and 15 user-defined key-value pairs (tags).
- **AWS Plugin on Panorama**—The Panorama plugin for AWS allows you to connect Panorama to your AWS VPC on the public cloud and retrieve the IP address-to-tag mapping for your virtual machines. Panorama then registers the VM information to the managed Palo Alto Networks® firewall(s) that you have configured for notification. With the plugin, Panorama can retrieve a total of 32 tags for each virtual machine, 15 predefined tags and up to 17 user-defined tags.

The maximum length of the tag-value (name and value included) must be 116 characters or less. If a tag is longer than 116 characters, Panorama does not retrieve the tag and register it on the firewalls.

Attributes Monitored on the AWS-VPC		VM Information Source on the Firewall	AWS Plugin on Panorama
AMI ID	ImageId.<ImageId string>	Yes	Yes
Architecture	Architecture.<Architecture string>	Yes	No
Availability Zone	AvailabilityZone.<string>	Yes	Yes
Guest OS	GuestOS.<guest OS name>	Yes	No
IAM Instance Profile	Iam-instance-profile.<instanceProfileArn>	No	Yes
Instance ID	InstanceId.<InstanceId string>	Yes	No
Instance State	InstanceState.<instance state>	Yes	No
Instance Type	InstanceType.<instance type>	Yes	No
Key Name	KeyName.<KeyName string>	Yes	Yes
Owner ID The value for this attribute is fetched from the ENI.	Account-number.<OwnerId>	No	Yes

 ©2024 Palo Alto Networks, Inc.

Attributes Monitored on the AWS-VPC		VM Information Source on the Firewall	AWS Plugin on Panorama
Placement Tenancy, Group Name	Placement.Tenancy.<string> Placement.GroupName.<string>	Yes	Yes
Private DNS Name	PrivateDnsName.<Private DNS Name>	Yes	No
Public DNS Name	PublicDnsName.<Public DNS Name>	Yes	Yes
Subnet ID	SubnetID.<subnetID string>	Yes	Yes
Security Group ID	Sg-id.<sg-xxxx>	No	Yes
Security Group Name	Sg-name.<SecurityGroupName>	No	Yes
VPC ID	VpcId.<VpcId string>	Yes	Yes
Tag (key, value)	aws-tag.<key>.<value>	Yes; Up to a maximum of 18 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 18 tags are available for use on the firewalls.	Yes; Up to a maximum of 21 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 21 tags are available for use on Panorama and the firewalls.
Load Balancer Name	Lb-name.<LoadBalancerName>	No	Yes
Endpoint ID	Endpoint-id.<EndpointID>	No	Yes

 ©2024 Palo Alto Networks, Inc.

IAM Permissions Required for Monitoring the AWS VPC

In order to enable VM Monitoring the user's AWS login credentials tied to the AWS Access Key and Secret Access Key must have permissions for the attributes listed above. These privileges allow the firewall to initiate API calls for monitoring the virtual machines in the AWS VPC.

The IAM policy associated with the user must either have global read-only access such as AmazonEC2ReadOnlyAccess, or must include individual permissions for all of the monitored attributes. The following IAM policy example lists the permissions for initiating the API actions for monitoring the resources in the AWS VPC:

```
{ “Version”: “2012-10-17",
  “Statement”: [
        {
            “Sid”: “VisualEditor0”,
            “Effect”: “Allow”,
            “Action”: [

 “elasticloadbalancing:DescribeLoadBalancerAttributes”,
                “elasticloadbalancing:DescribeLoadBalancers”,
                “elasticloadbalancing:DescribeTags”,
                “ec2:DescribeInstances”,
                “ec2:DescribeNetworkInterfaces”,
                “ec2:DescribeVpcs”,
                “ec2:DescribeVpcEndpoints”,
                “ec2:DescribeSubnets”
            ],
            “Resource”: “*”
        }
    ]
}
```

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on KVM

Kernel-based Virtual Machine (KVM) is an open-source virtualization module for servers running Linux distributions. The VM-Series firewall can be deployed on a Linux server that is running the KVM hypervisor.

This guide assumes that you have an existing IT infrastructure that uses Linux and have the foundation for using Linux/Linux tools. The instructions only pertain to deploying the VM-Series firewall on KVM.

- VM-Series on KVM— Requirements and Prerequisites
- Supported Deployments on KVM
- Install the VM-Series Firewall on KVM
- Performance Tuning of the VM-Series for KVM

VM-Series on KVM—Requirements and Prerequisites

- Options for Attaching the VM-Series on the Network
- Prerequisites for VM-Series on KVM

Table 5: VM-Series on KVM System Requirements

Requirements	Description
Hardware Resources	See VM-Series System Requirements for the minimum hardware requirements for your VM-Series model.
Software Versions	See the supported KVM software versions in the Compatibility Matrix.
SR-IOV Drivers	See PacketMMAP Driver Versions drivers in the Compatibility Matrix.
DPDK Drivers	See DPDK Driver Versions in the Compatibility Matrix. If you use one of the supported NIC drivers on VM-Series on KVM, DPDK is enabled by default.
Network Interfaces —Network Interface Cards and Software Bridges	The VM-Series on KVM supports a total of 25 interfaces — 1 management interface and a maximum of 24 network interfaces for data traffic. VM-Series deployed on KVM supports software-based virtual switches such as the Linux bridge or the Open vSwitch bridge, and direct connectivity to PCI passthrough or an SR-IOV capable adapter. If you plan to establish connectivity using PCI-passthrough or SR-IOV, you cannot configure a vSwitch on the physical port used for SR-IOV or PCI-passthrough. To communicate with the host and other virtual machines on the network, the VM-Series firewall must have exclusive access to the physical port and associated virtual functions (VFs) on that interface. • On the Linux bridge and OVS, the e1000 and Virtio drivers are supported; the default driver rtl8139 is not supported.

 ©2024 Palo Alto Networks, Inc.

Requirements	Description

	• For PCI passthrough/SR-IOV support, the VM-Series firewall has been tested for the following network cards:
	  • Intel 82576 based 1G NIC: SR-IOV support on all supported Linux distributions; PCI-passthrough support.
	  • Intel 82599 based 10G NIC: SR-IOV support on all supported Linux distributions; PCI-passthrough support.
	  • Intel X710 10G NIC: SR-IOV support on all supported Linux distributions; PCI-passthrough support
	  • Intel X722 10G NIC: SR-IOV support on all supported Linux distributions; PCI-passthrough support
	  • Broadcom 57112 and 578xx based 10G NIC: SR-IOV support on all supported Linux distributions; No PCI-passthrough support.
	  • Mellanox ConnectX5 10G/25G/50G/100G NIC: SR-IOV support on all supported Linux distributions.
	  • Refer to PacketMMAP Driver Versions in the Compatibility Matrix
	
	SR-IOV capable interfaces assigned to the VM-Series firewall, must be configured as Layer 3 interfaces or as HA interfaces.

Options for Attaching the VM-Series on the Network

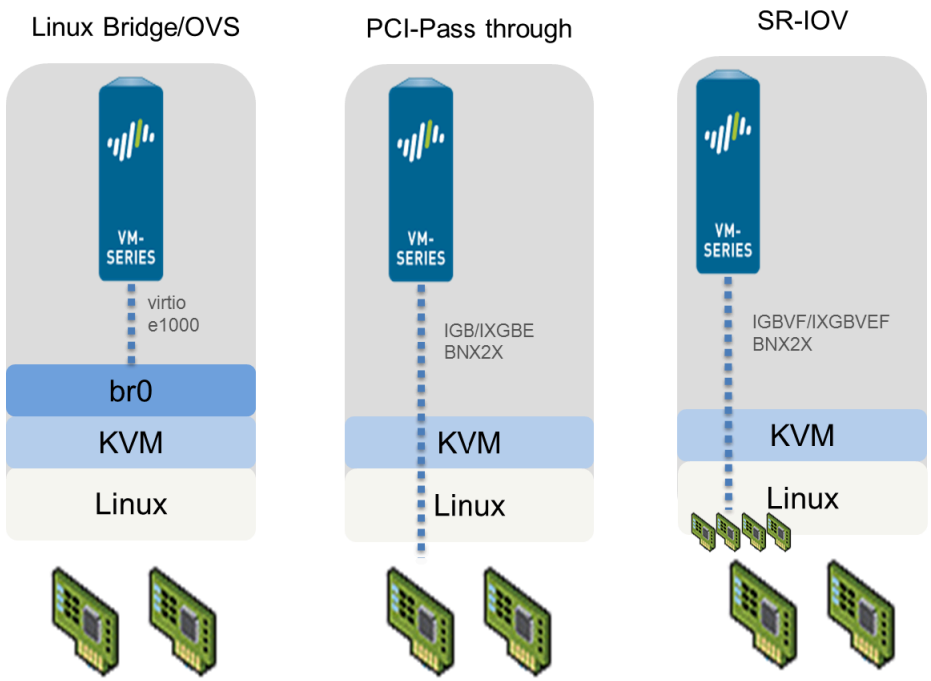

 ©2024 Palo Alto Networks, Inc.

- With a Linux bridge or OVS, data traffic uses the software bridge to connect guests on the same host. For external connectivity, data traffic uses the physical interface to which the bridge is attached.
- With PCI passthrough, data traffic is passed directly between the guest and the physical interface to which it is attached. When the interface is attached to a guest, it is not available to the host or to other guests on the host.
- With SR-IOV, data traffic is passed directly between the guest and the virtual function to which it is attached.

Prerequisites for VM-Series on KVM

Before you install the VM-Series firewall on the Linux server, review the following sections:

- Prepare the Linux Server
- Prepare to Deploy the VM-Series Firewall

Prepare the Linux Server

Before you install the VM-Series firewall on KVM, verify that you have a working Linux environment and that your networking infrastructure supports the connectivity a that your chosen deployment requires.

- Verify Linux Support
- Verify the Networking Infrastructure
- Install Mellanox Software Tools
- Enable Virtual Functions for Mellanox CX5 NICs on the VM-Series Firewall on KVM
- Verify the Host Configuration

Verify Linux Support

Verify that you have the correct environment to support your installation.

- ☐ Check the Linux distribution version. For a list of supported versions, see VM-Series for KVM in the Compatibility Matrix.
- ☐ Verify that you have installed and configured KVM tools and packages that are required for creating and managing virtual machines, such as Libvirt.
- ☐ If you want to use a SCSI disk controller to access the disk to which the VM-Series firewall stores data, you must use virsh to attach the virtio-scsi controller to the VM-Series firewall. You can then edit the XML template of the VM-Series firewall to enable the use of the virtio-scsi controller. For instructions, see Enable the Use of a SCSI Controller.

KVM on Ubuntu 12.04 does not support the virtio-scsi controller.

Verify the Networking Infrastructure

Verify that you have set up the networking infrastructure for steering traffic between the guests and the VM-Series firewall and ensure you have connectivity to an external server or the Internet. The VM-Series firewall can connect using a Linux bridge, the Open vSwitch, PCI passthrough, or SR-IOV capable network card.

 ©2024 Palo Alto Networks, Inc.

- Make sure that the link state for each interface you plan to use is Up—sometimes you have to manually bring up the interface.
- If using a Linux bridge or OVS, verify that you have set up the bridges required to send/receive traffic to/from the firewall. If not, create bridge(s) and verify that they are up before you begin installing the firewall.
- If using SR-IOV or PCI-passthrough, verify the PCI ID of all the interfaces. To view the list, use the following command:

  `Virsh nodedev-list –tree`

  See Verify PCI-ID for Ordering of Network Interfaces on the VM-Series Firewall.
- If using SR-IOV or PCI-passthrough, verify that the virtualization extensions (VT-d/IOMMU) are enabled in the BIOS. For example, to enable IOMMU, `intel_iommu=on` must be defined in `/etc/grub.conf`. Refer to the documentation provided by your system vendor for instructions.
- If using PCI-passthrough, ensure that the VM-Series firewall has exclusive access to the interface(s) that you plan to attach to it.

  To allow exclusive access, you must manually detach the interface(s) from the Linux server.

  `Virsh nodedev-detach <pci id of interface>`

  For example:

  `Virsh nodedev-detach pci_0000_07_10_0`

  In some cases, you might need to edit `/etc/libvirt/qemu.conf` and uncomment `relaxed_acs_check = 1`.
- If using SR-IOV, verify that the virtual function capability is enabled for each port that you plan to use on the network card. With SR-IOV, a single Ethernet port (physical function) can be split into multiple virtual functions. A guest can be mapped to one or more virtual functions.

  Enable virtual functions as follows:

  1. Create a new file in this location: `/etc/modprobe.d/`
  2. Use vi to edit the file to make the functions persistent:

     `vim /etc/modprobe.d/igb.conf`
  3. Enable the number of number of virtual functions required:

     `options igb max_vfs=4`

     In the above example, after you save the changes and reboot the Linux server, each interface (or physical function) will have 4 virtual functions.

  Refer to the documentation provided by your network vendor for details on the actual number of virtual functions supported, and instructions to enable virtual functions.

Install Mellanox Software Tools

If you are using a Mellanox CX5 card, install the Mellanox software tools on the host. Before installing, verify Linux support and your networking infrastructure.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | From the host, download the package for Mellanox OpenFabric Enterprise Distribution for Linux (MLNX_OFED) for your OS version from the following link:

https://www.mellanox.com/products/infiniband-drivers/linux/mlnx_ofed

STEP 2 | Run the installation command:

```
mlnxofedinstall
```

If you have all the prerequisite packages installed, the above command installs all the MLNX_OFED packages. Continue to Step 3.

If your environment doesn't have the required packages, the installer lists all the packages that you must install. After you install the packages, rerun the installation command and continue to Step 3.

STEP 3 | Reboot the host.

STEP 4 | Check the status of the Mellanox software tools.

```
# mst status
MST modules:
------------------------------------------------------------------
  MST PCI module is not loaded
  MST PCI configuration module loaded

MST devices:
------------------------------------------------------------------
/dev/mst/mt4121_pciconf0     - PCI configuration cycles access.
        domain:bus:dev.fn=0000:3b:00.0 addr.reg=88 data.reg=92
        Chip revision is: 00
```

STEP 5 | Ensure Mellanox is updated on the PCI list:

```
# lspci | grep Mellanox
3b:00.0 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex]
3b:00.1 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex]
```

Enable Virtual Functions for Mellanox CX5 NICs on the VM-Series Firewall on KVM

Install the Mellanox software tools before you enable virtual functions on Mellanox Cx5 NICs.

STEP 1 | Ensure Mellanox Software Tools (mst) are started.

STEP 2 | Enable the number of number of virtual functions required. For example:

```
mlxconfig -d /dev/mst/mt4121_pciconf0 set SRIOV_EN=1 NUM_OF_VFS=4
```

After you save the changes and reboot the Linux server, each interface (or physical function) in the above example will have 4 virtual functions. Refer to the documentation provided by

 ©2024 Palo Alto Networks, Inc.

your network vendor for details on the actual number of virtual functions supported, and the instructions to enable virtual functions.

You might see the following error message the first time you enable virtual functions on Mellanox Cx5 NICs:

```
[ 1429.841162] mlx5_core 0000:3b:00.1:
 mlx5_port_module_event:1025:(pid 0): Port module
 event[error]: module 1, Cable error, One or more network
 ports have been powered down due to insufficient/
unadvertised power on the PCIe slot. Please refer to the
 card's user manual for power specifications or contact
 Mellanox support
```

To resolve the issue, enter the following command sequence on the Linux server:

```
# mlxconfig -d <dev> set ADVANCED_POWER_SETTINGS=1
# mlxconfig -d <dev> set DISABLE_SLOT_POWER_LIMITER=1
# reboot
```

STEP 3 | Check the status of the virtual functions.

```
# cat  /sys/class/net/enp59s0f1/device/sriov_numvfs
```

(Optional) If the virtual functions are not set correctly (the status is 0 or empty), run the following command:

```
# echo 4 > /sys/class/net/enp59s0f1/device/sriov_numvfs
```

STEP 4 | List the PCI devices to accurately match the number of virtual functions loaded on the respective physical function for Mellanox:

```
# lspci | grep Mellanox
3b:00.0 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex]
3b:00.1 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex]
3b:00.2 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
3b:00.3 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
3b:00.4 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
3b:00.5 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
3b:00.6 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
3b:00.7 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
```

 ©2024 Palo Alto Networks, Inc.

```
3b:01.0 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
3b:01.1 Ethernet controller: Mellanox Technologies MT28800 Family
 [ConnectX-5 Ex Virtual Function]
```

Verify the Host Configuration

Configure the host for maximum VM-Series performance. Refer to Performance Tuning of the VM-Series for KVM for information about configuring each option below.

- **Enable DPDK**. DPDK allows the host to process packets faster by bypassing the Linux kernel. Instead, interactions with the NIC are performed using drivers and the DPDK libraries. Open vSwitch is required to use DPDK with the VM-Series firewall.
- **Enable SR-IOV**. Single root I/O virtualization (SR-IOV) allows a single PCIe physical device under a single root port to appear to be multiple separate physical devices to the hypervisor or guest.
- **Enable multi-queue support for NICs**. Multi-queue virtio-net allows network performance to scale with the number of vCPUs and allows for parallel packet processing by creating multiple TX and RX queues.
- **Isolate CPU resources in a NUMA Node**. You can improve performance of VM-Series on KVM by isolating the CPU resources of the guest VM to a single non-uniform memory access (NUMA) node.

Prepare to Deploy the VM-Series Firewall

- Purchase the VM-Series model and register the authorization code on the Palo Alto Networks Customer Support web site. See Create a Support Account and Register the VM-Series Firewall.
- Obtain the qcow2 image and save it on the Linux server. As a best practice, copy the image to the folder: /var/lib/libvirt/qemu/images.

  If you plan to deploy more than one instance of the VM-Series firewall, make the required number of copies of the image. Because each instance of the VM-Series firewall maintains a link with the .qcow2 image that was used to deploy the firewall, to prevent any data corruption issues ensure that each image is independent and is used by a single instance of the firewall.

 ©2024 Palo Alto Networks, Inc.

Supported Deployments on KVM

You can deploy a single instance of the VM-Series firewall per Linux host (single tenant) or multiple instances of the VM-Series firewalls on a Linux host. The VM-Series firewall can be deployed with virtual wire, Layer 2, or Layer 3 interfaces. If you plan on using SR-IOV capable interfaces on the VM-Series firewall, you can only configure the interfaces as Layer 3 interfaces.

- Secure Traffic on a Single Host
- Secure Traffic Across Linux hosts

Secure Traffic on a Single Host

To secure east west traffic across guests on a Linux server, the VM-Series firewall can be deployed with virtual wire, Layer 2, or Layer 3 interfaces. The illustration below shows the firewall with Layer 3 interfaces, where the firewall and the other guests on the server are connected using Linux bridges. In this deployment, all traffic between the web servers and the database servers is routed through the firewall; traffic across the database servers only or across the web servers only is processed by the bridge and is not routed through the firewall.

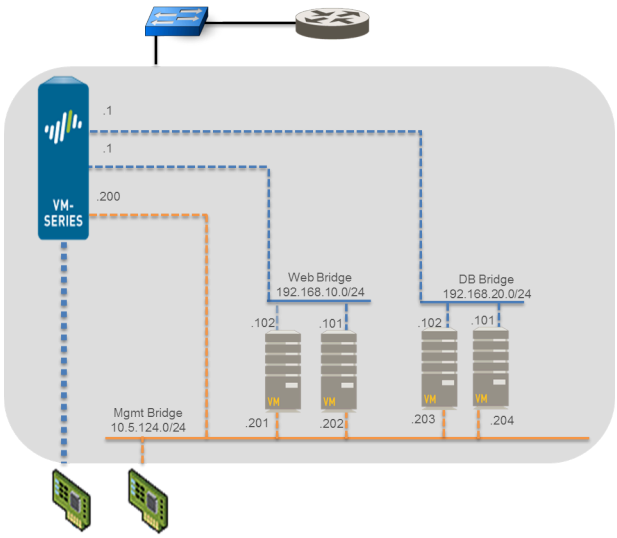

Secure Traffic Across Linux hosts

To secure your workloads, more than one instance of the VM-Series firewalls can be deployed on a Linux host. If, for example, you want to isolate traffic for separate departments or customers, you can use VLAN tags

to logically isolate network traffic and route it to the appropriate VM-Series firewall. In the following example, one Linux host hosts the VM-Series firewalls for two customers, Customer A and Customer B, and the workload for Customer B is spread across two servers. In order to isolate traffic and direct it to the VM-Series firewall configured for each customer, VLANs are used.

 ©2024 Palo Alto Networks, Inc.

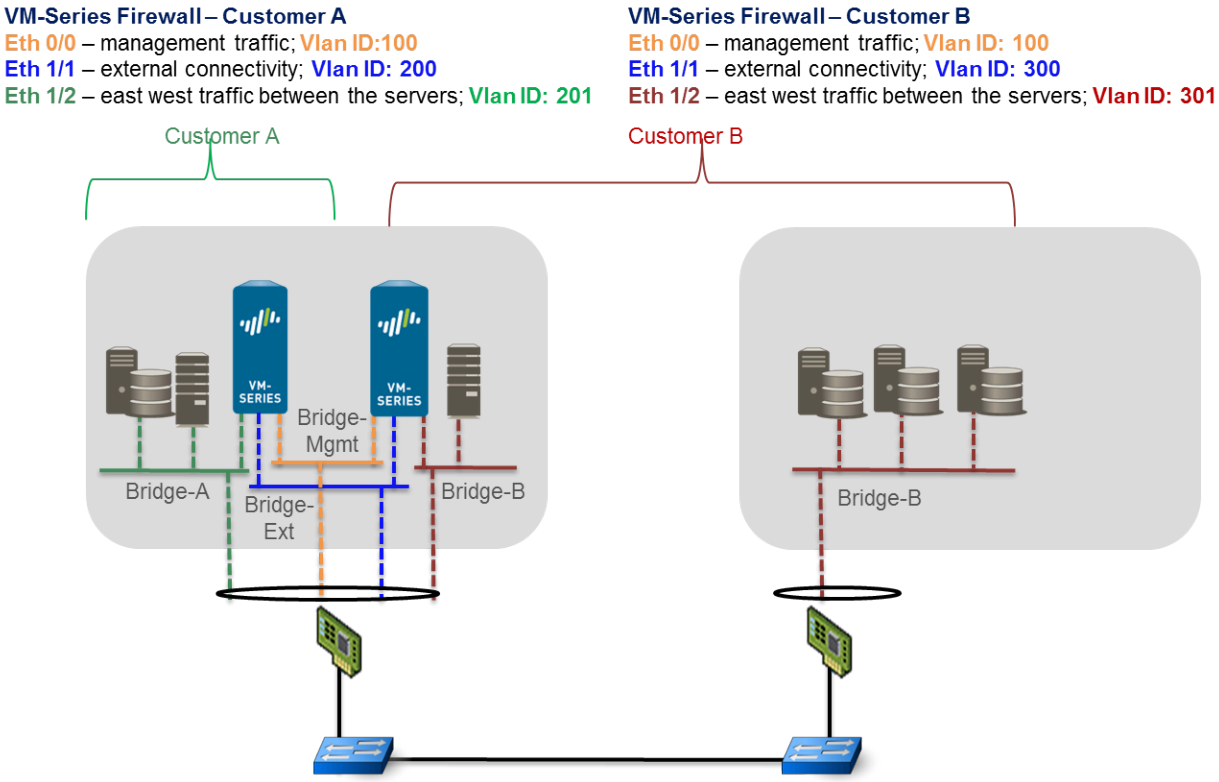

In another variation of this deployment, a pair of VM-Series firewalls are deployed in a high availability set up. The VM-Series firewalls in the following illustration are deployed on a Linux server with SR-IOV capable adapters. With SR-IOV, a single Ethernet port (physical function) can be split into multiple virtual functions. Each virtual function attached to the VM-Series firewall is configured as a Layer 3 interface. The active peer in the HA pair secures traffic that is routed to it from guests that are deployed on a different Linux server.

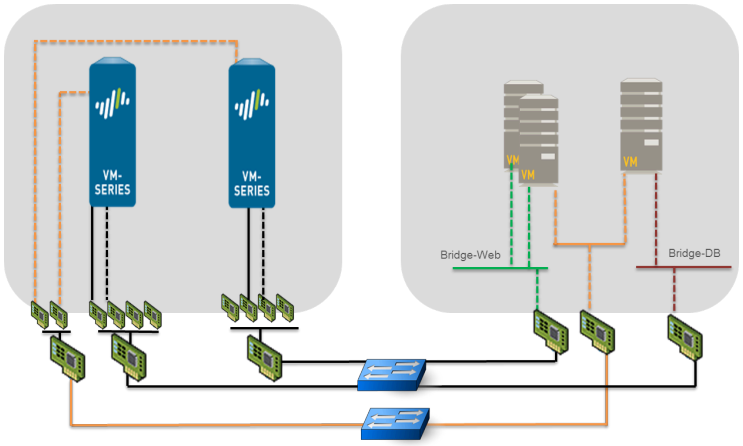

 ©2024 Palo Alto Networks, Inc.

Install the VM-Series Firewall on KVM

The libvirt API that is used to manage KVM includes a host of tools that allow you to create and manage virtual machines. To install the VM-Series firewall on KVM you can use any of the following methods.

- virt-manager—Deploy the VM-Series using the virt-manager virtual machine manager. Virt-manager provides a convenient wizard to help you through the installation process.
- virsh—Deploy the VM-Series using the KVM command line. Create an XML file that defines the virtual machine instance and bootstrap XML file that defines the initial configuration settings of the firewall. Then install the firewall by mounting an ISO image as a CD-ROM.
- virt-install—Another option to deploy the VM-Series firewall using the KVM command line. Use this option to create the definition for the VM-Series firewall and install it.

This document provides steps for installing the VM-Series firewall on KVM using virt-manager and virsh.

- Install the VM-Series Firewall Using Virt-Manager
- Install the VM-Series Firewall Using an ISO
- Use the VM-Series CLI to Swap the Management Interface on KVM

Install the VM-Series Firewall Using Virt-Manager

Use the following procedure uses virt-manager to install the VM-Series firewall on a server running KVM on RHEL.

- Provision the VM-Series Firewall on a KVM Host
- Perform Initial Configuration of the VM-Series Firewall on KVM

Provision the VM-Series Firewall on a KVM Host

Use the following instructions to provision the KVM host for the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create a new virtual machine and add the VM-Series Firewall for KVM image to virt-mgr.

1. On the Virt-manager, select **Create a new virtual machine**.
2. Add a descriptive **Name** for the VM-Series firewall.

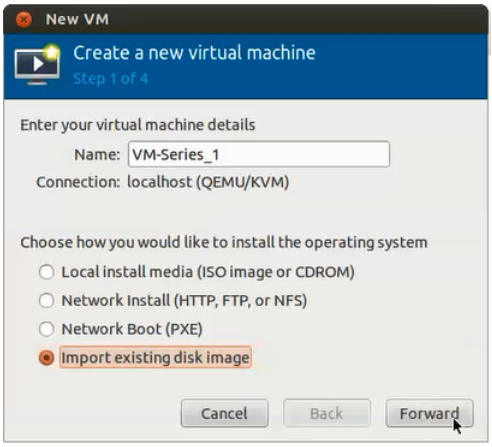

3. Select **Import existing disk image**, browse to the image, and set the **OS Type**: Linux and **Version**: Red Hat Enterprise Linux 6.

If you prefer, you can leave the OS Type and Version as Generic.

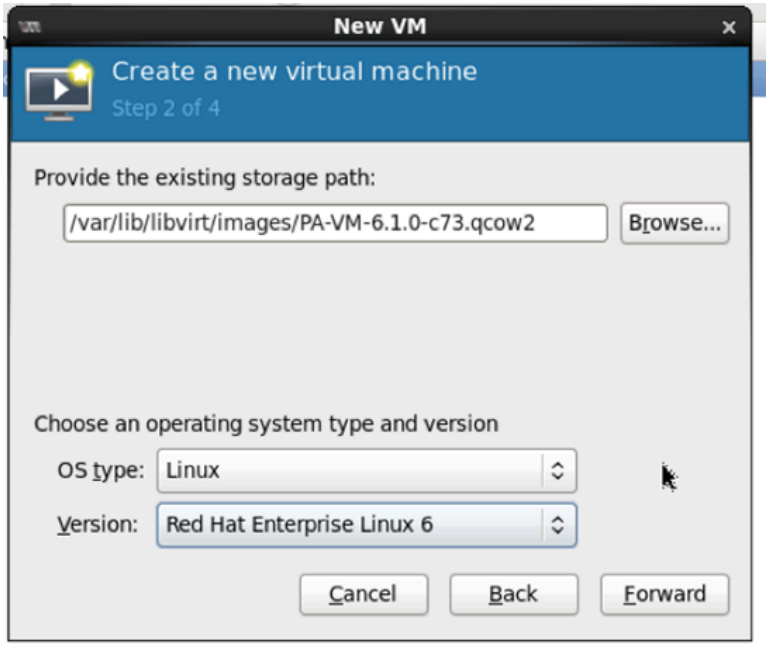

4. To add network adapters for the data interfaces:

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure the memory and CPU settings.

1. Set the **Memory** to the minimum memory based on the VM-Series System Requirements of your VM-Series model.
2. Set **CPU** to the minimum CPUs based on the VM-Series System Requirements of your VM-Series model.

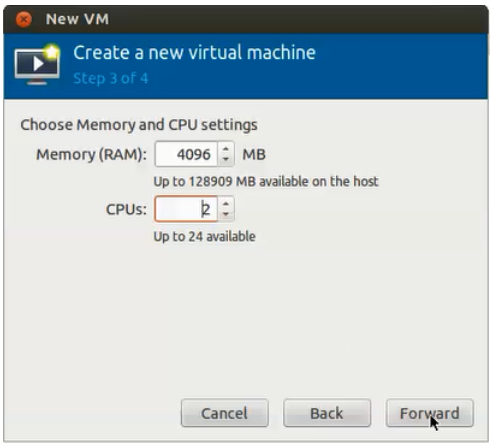

STEP 3 | Enable configuration customization and select the management interface bridge.

1. Select **Customize configuration before install**.
2. Under Advanced options, select the bridge for the management interface, and accept the default settings.

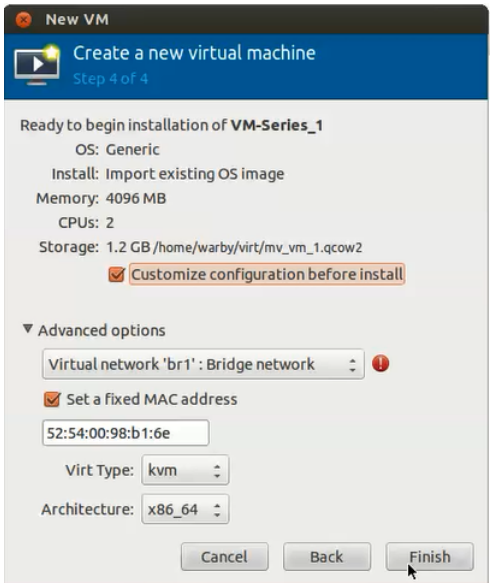

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Configure virtual disk settings.

1. Select **Disk**, expand Advanced options and select **Storage format** — qcow2; **Disk Bus**— Virtio or IDE, based on your set up.

   If you want to use a SCSI disk bus, see Enable the Use of a SCSI Controller.

2. Expand Performance options, and set **Cache mode** to **writethrough**. This setting improves installation time and execution speed on the VM-Series firewall.

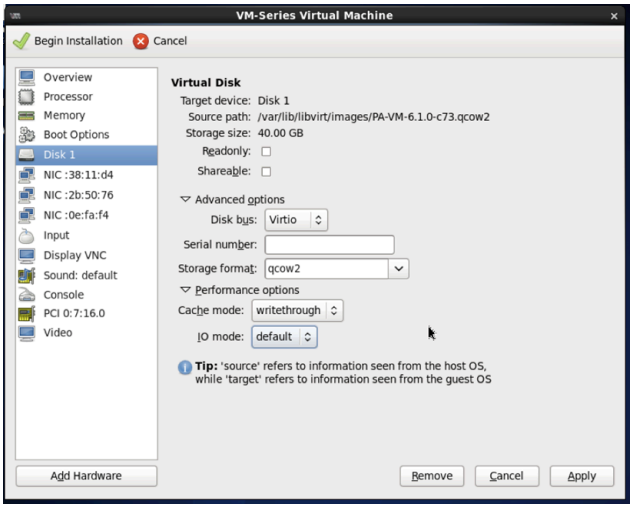

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Configure network adapters.

1. Select **Add Hardware** > **Network** if you are using a software bridge such as the Linux bridge or the Open vSwitch.
2. For **Host Device**, enter the name of the bridge or select it from the drop down list.
3. To specify the driver, set **Device Model** to e-1000 or virtio. These are the only supported virtual interface types.

4. Select **Add Hardware** > **PCI Host Device** for PCI-passthrough or an SR-IOV capable device.

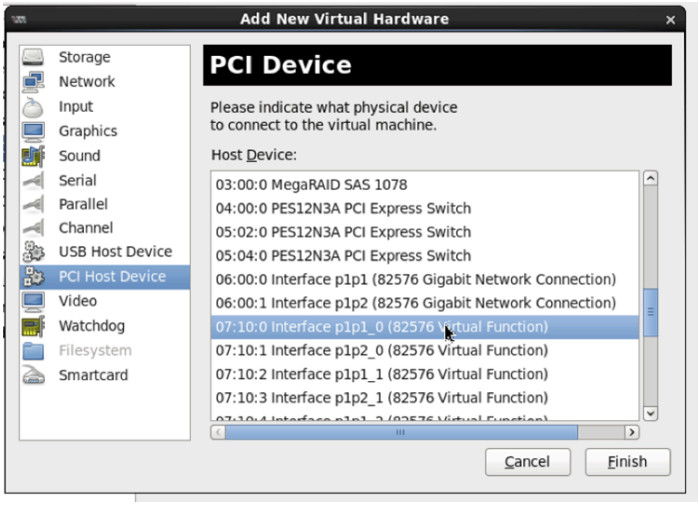

5. In the **Host Device** list, select the interface on the card or the virtual function.
6. Click **Apply** or **Finish**.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Click **Begin Installation** Begin Installation. Wait 5-7 minutes for the installation to complete.

By default, the XML template for the VM-Series firewall is created and stored at etc/libvirt/qemu.

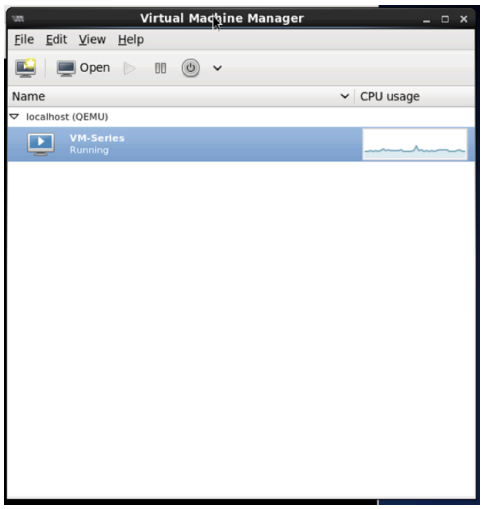

STEP 7 | (Optional) Bootstrap the VM-Series firewall

If you are using bootstrapping to perform the configuration of your VM-Series firewall on KVM, refer to Bootstrap the VM-Series Firewall on KVM. For more information about bootstrapping, see Bootstrap the VM-Series Firewall.

STEP 8 | Configure the network access settings for the management interface.

1. Open a connection to the console.
2. Log into the firewall with username/password: admin/admin.
3. Enter configuration mode with the following command:

```
configure
```

4. Use the following commands to configure the management interface:
   1.
```
set deviceconfig system type static
```
   2.
```
set deviceconfig system ip-address <Firewall-IP>
 netmask <netmask> default-gateway <gateway-IP> dns-setting
 servers primary <DNS-IP>
```

      where <*Firewall-IP*> is the IP address you want to assign to the management interface, <*netmask*> is the subnet mask, <*gateway-IP*> is the IP address of the network gateway, and <*DNS-IP*> is the IP address of the DNS server.
   3.
```
commit
```

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Verify which ports on the host are mapped to the interfaces on the VM-Series firewall. In order to verify the order of interfaces on the Linux host, see Verify PCI-ID for Ordering of Network Interfaces on the VM-Series Firewall.

To make sure that traffic is handled by the correct interface, use the following command to identify which ports on the host are mapped to the ports on the VM-Series firewall.

```
admin@PAN-VM> debug show vm-series interfaces all
Phoenix_interface    Base-OS_port       Base-OS_MAC PCI-ID
 mgt                        eth0          52:54:00:d7:91:52
 0000:00:03.0
 Ethernet1/1             eth1          52:54:00:fe:8c:80
 0000:00:06.0
 Ethernet1/2             eth2          0e:c6:6b:b4:72:06
 0000:00:07.0
 Ethernet1/3             eth3          06:1b:a5:7e:a5:78
 0000:00:08.0
 Ethernet1/4             eth4          26:a9:26:54:27:a1
 0000:00:09.0
 Ethernet1/5             eth5          52:54:00:f4:62:13
 0000:00:10.0
```

STEP 10 | Access the web interface of the VM-Series firewall and configure the interfaces and define security rules and NAT rules to safely enable the applications that you want to secure.

Refer to the PAN-OS Administrator's Guide.

Perform Initial Configuration of the VM-Series Firewall on KVM

Use the virtual appliance console on the KVM server to set up network access to the VM-Series firewall. By default, the VM-Series firewall uses DHCP to obtain an IP address for the management interface. However, you can assign a static IP address. After completing the initial configuration, access the web interface to complete further configurations tasks. If you have Panorama for central management, refer to the Panorama Administrator's Guide for more information on managing the device using Panorama.

If you are using bootstrapping to perform the configuration of your VM-Series firewall on KVM, refer to Bootstrap the VM-Series Firewall on KVM.

For general information about bootstrapping, see Bootstrap the VM-Series Firewall.

STEP 1 | Gather the required information from your network administrator.

- IP address for MGT port
- Netmask
- Default gateway
- DNS server IP address

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Access the console of the VM-Series firewall.

1. Select the **Console** tab on the KVM server for the VM-Series firewall, or right-click the VM-Series firewall and select **Open Console**.
2. Press enter to access the login screen.
3. Enter the default username/password (admin/admin) to log in.
4. Enter **`configure`** to switch to configuration mode.

STEP 3 | Configure the network access settings for the management interface.

Enter the following commands:

```
set deviceconfig system type static
```

```
set deviceconfig system ip-address <Firewall-IP> netmask <netmask>
default-gateway <gateway-IP> dns-setting servers primary <DNS-IP>
```

STEP 4 | Commit your changes and exit the configuration mode.

Enter **`commit`**.

Enter **`exit`**.

Install the VM-Series Firewall Using an ISO

Manually create the XML definition of the VM-Series firewall, then use virsh to import the definition as an ISO. Virsh is the most powerful tool that allows for full administration of the virtual machine.

- Use an ISO File to Deploy the VM-Series Firewall
- Sample XML file for the VM-Series Firewall

Use an ISO File to Deploy the VM-Series Firewall

If you want to pass a script to the VM-Series firewall at boot time, you can mount a CD-ROM with an ISO file. The ISO file allows you to define a bootstrap XML file that includes the initial configuration parameters for the management port of the firewall. The VM-Series firewall on first boot checks for the **bootstrap-networkconfig.xml** file, and uses the values defined in it.

If a single error is encountered in parsing the bootstrap file, the VM-Series firewall will reject all the configuration in this file and boot with default values.

STEP 1 | Create the XML file and define it as a virtual machine instance.

For a sample file, see Sample XML file for the VM-Series Firewall.

In this example, the VM-Series firewall is called PAN_Firewall_DC1.

For example:

```
user-PowerEdge-R510:~/kvm_script$ sudo vi /etc/libvirt/qemu/
PAN_Firewall_DC1.xml
```

 ©2024 Palo Alto Networks, Inc.

```
user-PowerEdge-R510:~/kvm_script$ sudo virsh define/etc/libvirt/
qemu/PAN_Firewall_DC1.xml
Domain PAN_Firewall_DC1_bootstp defined from /etc/libvirt/qemu/
PAN_Firewall_DC1.xml
user-PowerEdge-R510:~/kvm_script$ sudo virsh -q attach-interface
 PAN_Firewall_DC1_bootstp bridge br1 --model=virtio --persistent
user-PowerEdge-R510:~/kvm_script$ virsh list --all
 Id    Name                           State
----------------------------------------------
 -    PAN_Firewall_DC1_bootstp     shut off
```

STEP 2 | Create the bootstrap XML file.

You can define the initial configuration parameters in this file and name it bootstrap-networkconfig.

If you do not want to include a parameter, for example panorama-server-secondary. Delete the entire line from the file. If you leave the IP address field empty, the file will not be parsed successfully.

Use the following example as a template for the bootstrap-networkconfig file. The bootstrap-networkconfig file can include the following parameters only:

```
<vm-initcfg>
<hostname>VM_ABC_Company</hostname>
<ip-address>10.5.132.162</ip-address>
<netmask>255.255.254.0</netmask>
<default-gateway>10.5.132.1</default-gateway>
<dns-primary>10.44.2.10</dns-primary>
<dns-secondary>8.8.8.8</dns-secondary>
<panorama-server-primary>10.5.133.4</panorama-server-primary>
<panorama-server-secondary>10.5.133.5</panorama-server-secondary>
</vm-initcfg>
```

STEP 3 | Create the ISO file. In this example, we use mkisofs.

Save the ISO file in the images directory (/var/lib/libvirt/image) or the qemu directory (/etc/libvirt/qemu) to ensure that the firewall has read access to the ISO file.

For example:

```
# mkisofs -J -R -v -V "Bootstrap" -A "Bootstrap" -ldots -l -
allow-lowercase -allow-multidot -o <iso-filename> bootstrap-
networkconfig.xml
```

STEP 4 | Attach the ISO file to the CD-ROM.

For example:

```
# virsh -q attach-disk <vm-name> <iso-filename> sdc --type cdrom --
mode readonly –persistent\
```

 ©2024 Palo Alto Networks, Inc.

Sample XML file for the VM-Series Firewall

```
<?xml version="1.0"?>
<domain type="kvm">
<name>PAN_Firewall_DC1</name>
<memory>4194304</memory>
<currentMemory>4194304</currentMemory>
<vcpu placement="static">2</vcpu>
<os>
<type arch="x86_64">hvm</type>
<boot dev="hd"/>
</os>
<features>
<acpi/>
<apic/>
<pae/>
</features>
<clock offset="utc"/>
<on_poweroff>destroy</on_poweroff>
<on_reboot>restart</on_reboot>
<on_crash>restart</on_crash>
<devices>
<emulator>/usr/libexec/qemu-kvm</emulator>
<disk type="file" device="disk">
<driver type="qcow2" name="qemu"/>
<source file="/var/lib/libvirt/images/panos-kvm.qcow2"/>
<target dev="vda" bus="virtio"/>
</disk>
<controller type="usb" index="0"/>
<controller type="ide" index="0"/>
<controller type="scsi" index="0"/>
<serial type="pty">
<source path="/dev/pts/1"/>
<target port="0"/>
<alias name="serial0"/>
</serial>
<console type="pty" tty="/dev/pts/1">
<source path="/dev/pts/1"/>
<target type="serial" port="0"/>
<alias name="serial0"/>
</console>
<input type="mouse" bus="ps2"/>
<graphics type="vnc" port="5900" autoport="yes"/>
</devices>
</domain>
```

*To modify the number of vCPUs assigned on the VM-Series firewall, change the value **2** to **4** or **8** vCPUs in this line of the sample XML file:*

```
<vcpu placement="static">2</vcpu>
```

 ©2024 Palo Alto Networks, Inc.

Use the VM-Series CLI to Swap the Management Interface on KVM

By default, the VM-Series firewall assigns the first interface (eth0) as the management interface. However, in some deployments, the first interface must be pre-mapped to a public IP address. Therefore, the management interface must be assigned to a different interface. Assigning a public IP address to the management interface is a security risk.

Alternatively, you can enable management interface swap as part of the init-cfg.txt File Components *when bootstrapping.*

STEP 1 | Log in to the VM-Series firewall CLI and enter the following command:

```
set system setting mgmt-interface-swap enable yes
```

STEP 2 | Confirm that you want to swap the interface and use the eth1 dataplane interface as the management interface.

STEP 3 | Reboot the firewall for the swap to take effect. Use the following command:

```
request restart system
```

STEP 4 | Verify that the interfaces have been swapped. Use the following command:

```
debug show vm-series interfaces all
Phoenix_interface   Base-OS_port   Base-OS_MAC        PCI-ID
       Driver
mgt(interface-swap) eth0   0e:53:96:91:ef:29   0000:00:04.0
  ixgbevf
Ethernet1/1           eth1    0e:4d:84:5f:7f:4d       0000:00:03.0
  ixgbevf
```

 ©2024 Palo Alto Networks, Inc.

Enable the Use of a SCSI Controller

If you want the VM-Series firewall to use the disk bus type SCSI to access the virtual disk, use the following instructions to attach the virtio scsi controller to the firewall and then enable the use of the virtio-scsi controller.

KVM on Ubuntu 12.04 does not support the virtio-scsi controller; the virtio-scsi controller can only be enabled on the VM-Series firewall running on RHEL or CentOS.

This process requires virsh because Virt manager does not support the virtio-scsi controller.

STEP 1 | Create an XML file for the SCSI controller. In this example, it is called virt-scsi.xml.

```
[root@localhost~]# cat /root/virt-scsi.xml
<controller type='scsi' index='0' model='virtio-scsi'>
<address type='pci' domain='0x0000' bus='0x00'
 slot='0x0b'function='0x0'/>
</controller>
```

Make sure that the slot used for the virtio-scsi controller does not conflict with another device.

STEP 2 | Associate this controller with the XML template of the VM-Series firewall.

```
[root@localhost~]# virsh attach-device --config <VM-Series_name> /
root/virt-scsi.xml
Device attached successfully
```

STEP 3 | Enable the firewall to use the SCSI controller.

```
[root@localhost~]# virsh attach-disk <VM-Series_name>/var/lib/
libvirt/images/PA-VM-6.1.0-c73.qcow2
sda --cache none --persistent
Disk attached successfully
```

STEP 4 | Edit the XML template of the VM-Series firewall. In the XML template, you must change the target disk and the disk bus, used by the firewall.

By default, the XML template is stored at etc/libvirt/qemu.

```
<disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' cache='writeback'/>
      <source file='/var/lib/libvirt/images/PA-VM-7.0.0-c73.qcow2'/
>
      <target dev='sda' bus='scsi'/>
      <address type='drive' controller='0' bus='0' target='0'
 unit='0'/>
```

 ©2024 Palo Alto Networks, Inc.

```
</disk>
```

 ©2024 Palo Alto Networks, Inc.

Verify PCI-ID for Ordering of Network Interfaces on the VM-Series Firewall

Regardless of whether you use a virtual interfaces (Linux/OVS bridge) or PCI devices (PCI-passthrough or SR-IOV capable adapter) for connectivity to the VM-Series firewall, the VM-Series firewall treats the interface as a PCI device. The assignment of an interface on the VM-Series firewall is based on PCI-ID which is a value that combines the bus, device or slot, and function of the interface. The interfaces are ordered starting at the lowest PCI-ID, which means that the management interface (eth0) of the firewall is assigned to the interface with the lowest PCI-ID.

Let's say you assign four interfaces to the VM-Series firewall, three virtual interfaces of type virtio and e1000 and the fourth is a PCI device. To view the PCI-ID for each interface, enter the command **virsh dumpxml $ domain *<name of the VM-Series firewall>*** on the Linux host to view the list of interfaces attached to the VM-Series firewall. In the output, check for the following networking configuration:

```
<interface type='bridge'>
      <mac address='52:54:00:d7:91:52'/>
      <source bridge='mgmt-br'/>
      <model type='virtio'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x03'
 function='0x0'/>
    </interface>

   <interface type='bridge'>
      <mac address='52:54:00:f4:62:13'/>
      <source bridge='br8'/>
      <model type='e1000'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x10'
 function='0x0'/>
    </interface>

   <interface type='bridge'>
      <mac address='52:54:00:fe:8c:80'/>
      <source bridge='br8'/>
      <model type='e1000'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x06'
 function='0x0'/>
    </interface>

   <hostdev mode='subsystem' type='pci' managed='yes'>
      <source>
        <address domain='0x0000' bus='0x08' slot='0x10'
 function='0x1'/>
      </source>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x07'
 function='0x0'/>
    </hostdev>
```

In this case, the PCI-ID of each interface is as follows:

- First virtual interface PCI-ID is 00:03:00

 ©2024 Palo Alto Networks, Inc.

- Second virtual interface PCI-ID is 00:10:00
- Third virtual interface PCI-ID is 00:06:00
- Fourth interface PCI-ID is 00:07:00

Therefore, on the VM-Series firewall, the interface with PCI-ID of 00:03:00 is assigned as eth0 (management interface), the interface with PCI-ID 00:06:00 is assigned as eth1 (ethernet1/1), the interface with PCI-ID 00:07:00 is eth2 (ethernet1/2) and the interface with PCI-ID 00:10:00 is eth3 (ethernet1/3).

 ©2024 Palo Alto Networks, Inc.

Performance Tuning of the VM-Series for KVM

The VM-Series firewall for KVM is a high-performance appliance but may require tuning of the hypervisor to achieve the best results. This section describes some best practices and recommendations for facilitating the best performance of the VM-Series firewall.

By default, KVM uses a linux bridge for VM networking. However, the best performance in a virtual environment is realized with dedicated I/O interfaces (PCI passthrough or SR-IOV). If a virtual switch is required, use a performance-optimized virtual switch (such as Open vSwitch with DPDK).

- Install KVM and Open vSwitch on Ubuntu 16.04.1 LTS
- Enable Open vSwitch on KVM
- Integrate Open vSwitch with DPDK
- Enable SR-IOV on KVM
- Enable VLAN Access Mode with SR-IOV
- Enable Multi-Queue Support for NICs on KVM
- Isolate CPU Resources in a NUMA Node on KVM

Install KVM and Open vSwitch on Ubuntu 16.04.1 LTS

For ease of installation, Ubuntu 16.04.1 LTS is recommended for use as the KVM hypervisor platform.

STEP 1 | Install KVM and OVS.

1. Log in to the Ubuntu CLI.
2. Execute the following commands:

```
$ sudo apt-get install qemu-kvm libvirt-bin ubuntu-vm-builder
 bridge-utils
$ sudo apt-get install openvswitch-switch
```

STEP 2 | Check and compare the versions of relevant packages.

Execute the following commands:

```
$ virsh --version 1.3.1
$ libvirtd --version
libvirtd (libvirt) 1.3.1
$ /usr/bin/qemu-system-x86_64 --version
QEMU emulator version 2.5.0 (Debian
1:2.5+dfsg-5ubuntu10.6), Copyright (c) 2003-2008
Fabrice Bellard
$ ovs-vsctl --version
ovs-vsctl (Open vSwitch) 2.5.0
Compiled Mar 10 2016 14:16:49
DB Schema 7.12.1
```

 ©2024 Palo Alto Networks, Inc.

Enable Open vSwitch on KVM

Enable OVS by modifying the guest XML definition network settings.

Modify the guest XML definition as follows.

```
[...]
    <interface type='bridge'>
       <mac address='52:54:00:fb:00:01'/>
       <source bridge='ovsbr0'/>
       <virtualport type='openvswitch'/>
       <model type='virtio'/>
       <address type='pci' domain='0x0000' bus='0x00' slot='0x03'
 function='0x0'/>
   </interface>
[...]
```

Integrate Open vSwitch with DPDK

To integrate Open vSwitch (OVS) with DPDK, you must install the required components and then configure OVS. DPDK is enabled by default on the VM-Series firewall for KVM.

- Install QEMU, DPDK, and OVS on Ubuntu
- Configure OVS and DPDK on the Host
- Edit the VM-Series Firewall Configuration File

Install QEMU, DPDK, and OVS on Ubuntu

Before you can enable DPDK on OVS, you must install QEMU 2.5.0, DPDK 2.2.0, and OVS 2.5.1. Complete the following procedures to install the components.

STEP 1 | Log in to the KVM host CLI.

STEP 2 | Install QEMU 2.5.0 by executing the following commands:

```
apt-get install build-essential gcc pkg-config glib-2.0 libglib2.0-
dev libsdl1.2-dev
libaio-dev libcap-dev libattr1-dev libpixman-1-dev
apt-get build-dep qemu
apt-get install qemu-kvm libvirt-bin
wget http://wiki.qemu.org/download/qemu-2.5.0.tar.bz2
tar xjvf qemu-2.5.0.tar.bz2
cd qemu-2.5.0
./configure
make
make install
```

STEP 3 | Install dpdk-2.2.0.

1. Execute the following commands:

```
wget http://dpdk.org/browse/dpdk/snapshot/dpdk-2.2.0.tar.gz
```

 ©2024 Palo Alto Networks, Inc.

```
tar xzvf dpdk-2.2.0.tar.gz
cd dpdk-2.2.0
vi config/common_linuxapp
```

2. Change `CONFIG_RTE_APP_TEST=y` to **`CONFIG_RTE_APP_TEST=n`**
3. Change `CONFIG_RTE_BUILD_COMBINE_LIBS=n` to **`CONFIG_RTE_BUILD_COMBINE_LIBS=y`**
4. Execute the following command:

```
vi GNUmakefile
```

5. Change `ROOTDIRS-y := lib drivers app` to **`ROOTDIRS-y := lib drivers`**
6. Execute the following command:

```
make install T=x86_64-native-linuxapp-gcc
```

STEP 4 | Install OVS 2.5.1 by executing the following commands:

```
wget http://openvswitch.org/releases/openvswitch-2.5.1.tar.gz
tar xzvf openvswitch-2.5.1.tar.gz
cd openvswitch-2.5.1
./configure –with-dpdk=”/root/dpdk-2.2.0/x86_64-native-linuxapp-
gcc/”
make
make install
```

Configure OVS and DPDK on the Host

After installing the necessary components to support OVS and DPDK, you must configure the host to use OVS and DPDK.

STEP 1 | Log in to the KVM host CLI.

STEP 2 | If you are replacing or reconfiguring an existing OVS-DPDK setup, execute the following commands to reset any previous configuration. Repeat the command for each interface.

```
rm /usr/local/var/run/openvswitch/<interface-name>
```

STEP 3 | Configure initial huge pages for OVS.

```
echo 16384 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages
```

STEP 4 | Mount huge pages for QEMU:

```
mkdir /dev/hugepages
mkdir /dev/hugepages/libvirt
mkdir /dev/hugepages/libvirt/qemu
mount -t hugetlbfs hugetlbfs /dev/hugepages/libvirt/qemu
```

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Use the following command to kill any currently existing OVS daemon.

```
killall ovsdb-server ovs-vswitchd
```

STEP 6 | Create directories for the OVS daemon.

```
mkdir -p /usr/local/etc/openvswitch
mkdir -p /usr/local/var/run/openvswitch
```

STEP 7 | Clear old directories.

```
rm -f /var/run/openvswitch/vhost-user*
rm -f /usr/local/etc/openvswitch/conf.db
```

STEP 8 | Initialize the configuration database.

```
ovsdb-tool create /usr/local/etc/openvswitch/conf.db\
/usr/local/share/openvswitch/vswitch.ovsschema
```

STEP 9 | Create an OVS DB server.

```
ovsdb-server --remote=punix:/usr/local/var/run/openvswitch/db.sock
 \
--remote=db:Open_vSwitch,Open_vSwitch,manager_options \
--private-key=db:Open_vSwitch,SSL,private_key \
--certificate=db:Open_vSwitch,SSL,certificate \
--bootstrap-ca-cert=db:Open_vSwitch,SSL,ca_cert \
--pidfile --detach
```

STEP 10 | Initialize OVS.

```
ovs-vsctl --no-wait init
```

STEP 11 | Start the database server.

```
export DB_SOCK=/usr/local/var/run/openvswitch/db.sock
```

STEP 12 | Install the igb_uio module (network device driver) for DPDK.

```
cd ~/dpdk-2.2.0/x86_64-native-linuxapp-gcc/kmod
modprobe uio
insmod igb_uio.ko
cd ~/dpdk-2.2.0/tools/
```

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Enable DPDK on interfaces using PCI-ID or interface name.

```
./dpdk_nic_bind.py --bind=igb_uio <your first data interface>
./dpdk_nic_bind.py --bind=igb_uio <your second data interface>
```

STEP 14 | Start the OVS daemon in DPDK mode. You can change the number of cores for ovs-vswitchd. By changing -c 0x1 to -c 0x3, you can have two core run this daemon.

```
ovs-vswitchd --dpdk -c 0x3 -n 4 -- unix:$DB_SOCK --pidfile --
detach
echo 50000 > /sys/kernel/mm/hugepages/hugepages-2048kB/
nr_hugepages
```

STEP 15 | Create the OVS bridge and attach ports to the OVS bridge.

```
ovs-vsctl add-br ovs-br0 -- set bridge ovs-br0
 datapath_type=netdev
ovs-vsctl add-port ovs-br0 dpdk0 -- set Interface dpdk0 type=dpdk
ovs-vsctl add-br ovs-br1 -- set bridge ovs-br1
 datapath_type=netdev
ovs-vsctl add-port ovs-br1 dpdk1 -- set Interface dpdk1 type=dpdk
```

STEP 16 | Create DPDK vhost user ports for OVS.

```
ovs-vsctl add-port ovs-br0 vhost-user1 -- set Interface vhost-user1
 type=dpdkvhostuser
ovs-vsctl add-port ovs-br1 vhost-user2 -- set Interface vhost-user2
 type=dpdkvhostuser
```

STEP 17 | Set the number of hardware queues of the NIC used by the host.

```
ovs-vsctl set Open_vSwitch . other_config:n-dpdk-rxqs=8
ovs-vsctl set Open_vSwitch . other_config:n-dpdk-txqs=8
```

STEP 18 | Set the CPU mask used for OVS.

```
ovs-vsctl set Open_vSwitch . other_config:pmd-cpu-mask=0xffff
```

STEP 19 | Set the necessary permissions for DPDK vhost user ports. In the example below, 777 is used to give read, write, and executable permissions.

```
chmod 777 /usr/local/var/run/openvswitch/vhost-user1
chmod 777 /usr/local/var/run/openvswitch/vhost-user2
chmod 777 /dev/hugepages/libvirt/qemu
```

 ©2024 Palo Alto Networks, Inc.

Edit the VM-Series Firewall Configuration File

Edit the VM-Series firewall XML configuration file to support OVS and DPDK. You can access the XML configuration file or after deploying the VM-Series firewall. If you do this after deploying the firewall, be sure to shut down the firewall before making any changes. The values below are examples, your values for each parameter will vary based on your VM-Series model.

STEP 1 | Log in to the KVM host CLI.

STEP 2 | Edit the XML configuration file of your VM-Series firewall.

1. Open the XML config file using **virsh edit $<*your-vm-series-name*>**.
2. Sets the memory backing for the hugepage. Ensure that you provide enough memory to support the VM-Series firewall model you are deploying on the host. See VM-Series System Requirements for more information.

```
<memory unit='KiB'>12582912</memory>
  <currentMemory unit='KiB'>6291456</currentMemory>
  <memoryBacking>
    <hugepages/>
```

3. Set the necessary CPU flags for VM.

```
<cpu mode='host-model'>
```

4. Enable memory sharing between the VM and the host.

```
<numa>
      <cell id='0' cpus='0,2,4,6' memory='6291456' unit='KiB'
 memAccess='shared'/>
      <cell id='1' cpus='1,3,5,7' memory='6291456' unit='KiB'
 memAccess='shared'/>
    </numa>
```

5. Set the DPDK vhost user ports as the VM -series firewall's network interfaces. Additionally, set the number of virtio virtual queues provided to the VM-Series firewall by the host.

```
<interface type='vhostuser'>
      <mac address='52:54:00:36:83:70'/>
      <source type='unix' path='/usr/local/var/run/
openvswitch/vhost-user1' mode='client'/>
<model type='virtio'/>
<driver name='vhost' queues='8'/>
      <address type='pci' domain='0x0000' bus='0x00'
 slot='0x04' function='0x0'/>
    </interface>
    <interface type='vhostuser'>
      <mac address='52:54:00:30:d7:94'/>
      <source type='unix' path='/usr/local/var/run/
openvswitch/vhost-user2' mode='client'/>
<model type='virtio'/>
<driver name='vhost' qeueus='8'>
```

 ©2024 Palo Alto Networks, Inc.

```
      <address type='pci' domain='0x0000' bus='0x00'
 slot='0x05' function='0x0'/>
   </interface>
```

Enable SR-IOV on KVM

Single root I/O virtualization (SR-IOV) allows a single PCIe physical device under a single root port to appear to be multiple separate physical devices to the hypervisor or guest. To enable SR-IOV on a KVM guest, define a pool of virtual function (VF) devices associated with a physical NIC and automatically assign VF devices from the pool to PCI IDs.

For SR-IOV with Intel 10GB network interfaces (ixgbe driver), the driver version must be 4.2.5 or later to support multiple queues for each NIC interface. See the Compatibility Matrix for PacketMMAP and DPDK driver support by PAN-OS version.

STEP 1 | Define a network for a pool of VFs.

1. Generate an XML file with text similar to the following example. Change the value of `pf dev` to the `ethdev` corresponding to your SR-IOV device's physical function.

```
<network>
   <name>passthrough</name>
   <forward mode='hostdev' managed='yes'>
     <pf dev='eth3'/>
   </forward>
</network>
```

2. Save the XML file.
3. Execute the following commands:

```
$ virsh net-define <path to network XML file>
$ virsh net-autostart passthrough
$ virsh net-start passthrough
```

STEP 2 | To ensure that the VM-Series firewall boots in DPDK mode, edit the guest VM XML configuration on the KVM hypervisor to add the following:

```
<cpu mode='host-passthrough' check='none'/>
```

This ensures that the CPU flags are exposed.

To verify that the CPU flags are exposed on the VM:

```
cat /proc/cpuinfo
```

In the `flags` output for PAN-OS 10.0 or later with DPDK 18.11, you need AVX, or AES and SSE flags.

STEP 3 | After defining and starting the network, modify the guest XML definition to specify the network.

```
<interface type='network'>
    <source network='passthrough'>
```

 ©2024 Palo Alto Networks, Inc.

```
</interface>
```

When the guest starts, a VF is automatically assigned to the guest.

STEP 4 | Add the multicast MAC address to the host.

When SR-IOV is enabled, multicast traffic is filtered by the PF. This filtering causes applications that rely on multicast, such as OSPF, to fail. To prevent this filtering, you must manually add the multicast MAC address to the host using the following command:

```
#ip maddress add <multicast-mac> dev <interface-name>
```

Enable VLAN Access Mode with SR-IOV

The VM-Series firewalls on KVM can operate in VLAN access mode to support use cases where it is deployed as a virtual network function (VNF) that offers security-as-a-service in a multi-tenant cloud/data center environment. In VLAN access mode, each VNF has dedicated virtual network interfaces (VNIs) for each network and it sends and receives packets to/from SR-IOV virtual functions (VFs) without VLAN tags; you must enable this capability on the physical and virtual functions on the host hypervisor. When you, then enable VLAN access mode on the VM-Series firewall, the firewall can send and receive traffic without VLAN tags across all its dataplane interfaces. Additionally, if you configure QoS policies, the firewall can enforce QoS on the access interface and provide differentiated treatment of traffic in a multi-tenant deployment.

By default, the VM-Series firewall on KVM operates in VLAN trunk mode.

STEP 1 | On the host system, set up the physical and virtual function to operate in VLAN access mode.

`ip link set [inf_name] vf [vf_num] vlan [vlan_id]`.

For best performance on the VM-Series firewall, make sure to:

- *Enable CPU pinning. See* Isolate CPU Resources in a NUMA Node on KVM.
- *Disable Replay Protection, if you have configured IPSec Tunnels.*

  On the firewall web interface, select ***Network*** > ***IPSec Tunnels*** *select an IPSec tunnel, and click* ***General****, and select* ***Show Advanced Options*** *and clear* ***Enable Replay Protection****.*

STEP 2 | Access the CLI on the VM-Series firewall.

STEP 3 | Enable VLAN access mode.

`request plugins vm-series vlan-mode access-mode on`

`on` enables VLAN access mode; to use VLAN trunk mode, enter **`request plugins vm-series vlan-mode access-mode off`**.

STEP 4 | Reboot the firewall.

Enter**`request restart system`**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Verify the VLAN mode configuration.

```
show plugins vm-series vlan-mode
```

Enable Multi-Queue Support for NICs on KVM

Modify the guest XML definition to enable multi-queue virtio-net. Multi-queue virtio-net allows network performance to scale with the number of vCPUs and allows for parallel packet processing by creating multiple TX and RX queues.

Modify the guest XML definition. Insert a value from 1 to 256 for N to specify the number of queues. For the best results, match the number of queues with number of dataplane cores configured on the VM.

```
<interface type='network'>
      <source network='default'/>
      <model type='virtio'/>
      <driver name='vhost' queues='N'/>
</interface>
```

Isolate CPU Resources in a NUMA Node on KVM

You can improve performance of VM-Series on KVM by isolating the CPU resources of the guest VM to a single non-uniform memory access (NUMA) node. On KVM, you can view the NUMA topology **virsh**. The following example is from a two-node NUMA system:

STEP 1 | View the NUMA topology. In the example below, there are two NUMA nodes (sockets), each with a four-core CPU with hyperthreading enabled. All the even-numbered CPU IDs belong to one node and all the odd-numbered CPU IDs belong to the other node.

```
% virsh capabilities
<…>
  <topology>
    <cells num='2'>
      <cell id='0'>
        <memory unit='KiB'>33027228</memory>
        <pages unit='KiB' size='4'>8256807</pages>
        <pages unit='KiB' size='2048'>0</pages>
        <distances>
          <sibling id='0' value='10'/>
          <sibling id='1' value='20'/>
        </distances>
        <cpus num='8'>
          <cpu id='0' socket_id='1' core_id='0' siblings='0,8'/>
          <cpu id='2' socket_id='1' core_id='1' siblings='2,10'/>
          <cpu id='4' socket_id='1' core_id='2' siblings='4,12'/>
          <cpu id='6' socket_id='1' core_id='3' siblings='6,14'/>
          <cpu id='8' socket_id='1' core_id='0' siblings='0,8'/>
          <cpu id='10' socket_id='1' core_id='1' siblings='2,10'/>
          <cpu id='12' socket_id='1' core_id='2' siblings='4,12'/>
          <cpu id='14' socket_id='1' core_id='3' siblings='6,14'/>
        </cpus>
      </cell>
```

 ©2024 Palo Alto Networks, Inc.

```
      <cell id='1'>
        <memory unit='KiB'>32933812</memory>
        <pages unit='KiB' size='4'>8233453</pages>
        <pages unit='KiB' size='2048'>0</pages>
        <distances>
          <sibling id='0' value='20'/>
          <sibling id='1' value='10'/>
        </distances>
        <cpus num='8'>
          <cpu id='1' socket_id='0' core_id='0' siblings='1,9'/>
          <cpu id='3' socket_id='0' core_id='1' siblings='3,11'/>
          <cpu id='5' socket_id='0' core_id='2' siblings='5,13'/>
          <cpu id='7' socket_id='0' core_id='3' siblings='7,15'/>
          <cpu id='9' socket_id='0' core_id='0' siblings='1,9'/>
          <cpu id='11' socket_id='0' core_id='1' siblings='3,11'/>
          <cpu id='13' socket_id='0' core_id='2' siblings='5,13'/>
          <cpu id='15' socket_id='0' core_id='3' siblings='7,15'/>
       </cpus>
    </cell>
</cells>
```

STEP 2 | Pin vCPUs in a KVM guest to specific physical vCPUs, use the **cpuset** attribute in the guest xml definition. In this example, all 8 vCPUs are pinned to physical CPUs in the first NUMA node. If you do not wish to explicitly pin the vCPUs, you can omit the **cputune** block, in which case, all vCPUs will be pinned to the range of CPUs specified in cpuset, but will not be explicitly mapped.

```
<vcpu cpuset='0,2,4,6,8,10,12,14'>8</vcpu>
<cputune>
  <vcpupin vcpu='0' cpuset='0'/>
  <vcpupin vcpu='1' cpuset='2'/>
  <vcpupin vcpu='2' cpuset='4'/>
  <vcpupin vcpu='3' cpuset='6'/>
  <vcpupin vcpu='4' cpuset='8'/>
  <vcpupin vcpu='5' cpuset='10'/>
  <vcpupin vcpu='6' cpuset='12'/>
  <vcpupin vcpu='7' cpuset='14'/>
</cputune>
```

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on Hyper-V

The VM-Series firewall can be deployed on a server running Microsoft Hyper-V. Hyper-V is packaged as a standalone hypervisor or as an add-on/role for Windows Server.

- Supported Deployments on Hyper-V
- System Requirements on Hyper-V
- Linux Integration Services
- Install the VM-Series Firewall on Hyper-V

Supported Deployments on Hyper-V

You can deploy one or more instances of the VM-Series on hosts running Hyper-V. Where you place the VM-Series firewall depends on your network topology. VM-Series supports tap, virtual wire, Layer 2, and Layer 3 interface deployments.

- Secure Traffic on a Single Hyper-V Host
- Secure Traffic Across Multiple Hyper-V Hosts

Secure Traffic on a Single Hyper-V Host

The VM-Series firewall is deployed on a single Hyper-V host along with other guest VMs. In the example below, the VM-Series firewall has a Layer 3 interfaces and the VM-Series and other guest VMs are connected by Hyper-V vSwitches. All traffic between the web servers and database servers is routed through the firewall. Traffic across the database servers only or across the web servers only is processed by the external vSwitch and not routed through the firewall.

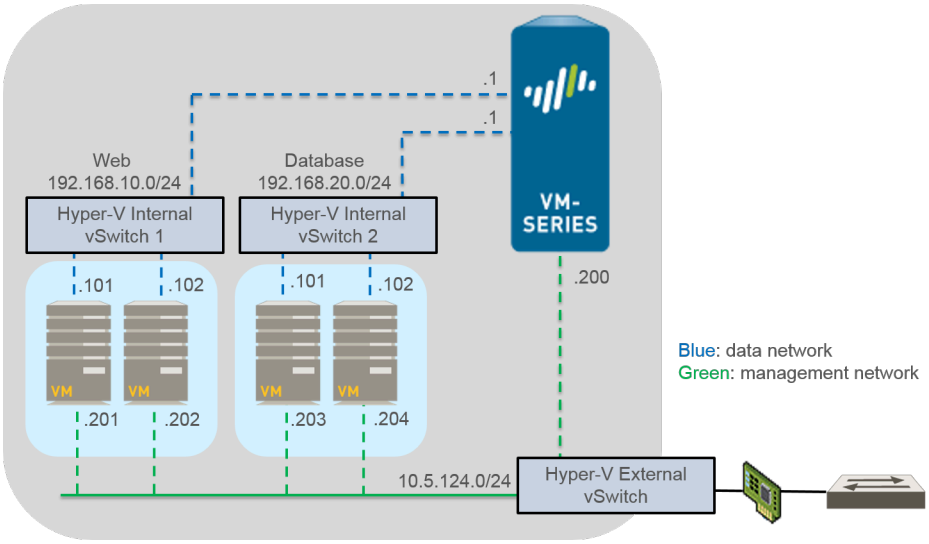

Secure Traffic Across Multiple Hyper-V Hosts

You can deploy your VM-Series firewall to secure the traffic of multiple Hyper-V hosts. In the example below, the VM-Series is deployed in Layer 2 mode protecting traffic to and from the guest VMs. A single VM-Series firewall protects traffic between four guest VMs spread across two Hyper-V hosts. VLAN tagging is used to logically isolate traffic and direct it to the firewall. Additionally, management traffic is decoupled from all other traffic by placing it on its own external vSwitch.

 ©2024 Palo Alto Networks, Inc.

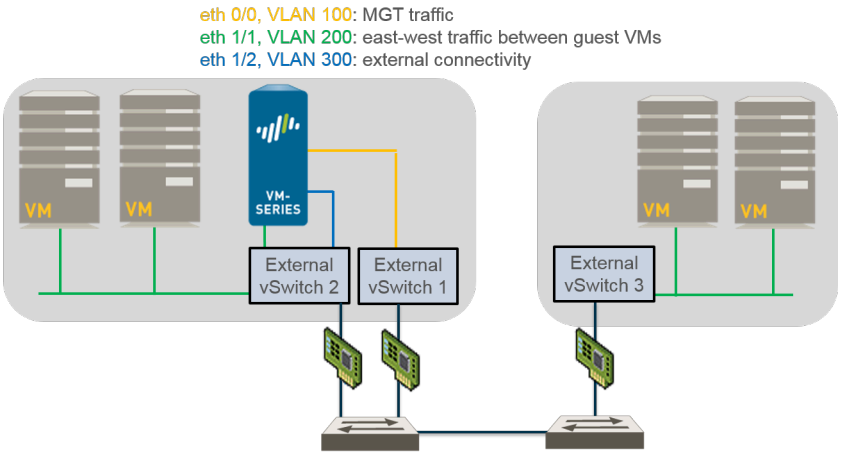

 ©2024 Palo Alto Networks, Inc.

System Requirements on Hyper-V

The VM-Series requires a minimum resource allocation on the Hyper-V host, so make sure to conform to the requirements listed below to ensure optimal performance.

- The host CPU must be a 64-bit x86-based Intel or AMD CPU with virtualization extension.
- See VM-Series System Requirements for the minimum hardware requirements for your VM-Series model.
- Minimum of two network adapters. The VM-Series firewall supports synthetic network adapters, which provide better performance than emulated network adapters. Hyper-V supports up to eight synthetic network adapters.
- Refer to the Compatibility matrix for the Windows Server versions supported.

  Hyper-V Server does not have a native graphical user interface; all configuration is done through PowerShell. However, you can use Hyper-V Manager running on a remote machine to manage the firewall.If you use the Hyper-V role add-on, you can manage the firewall using Hyper-V Manager or PowerShell.
- The VM-Series firewall does not support Legacy Network Adapter or SR-IOV/PCI-Passthrough.

 ©2024 Palo Alto Networks, Inc.

Linux Integration Services

Linux Integration Services (LIS) is a package of drivers and services that enhance the performance of Linux-based virtual machines on Hyper-V. The VM-Series firewall supports the following services to improve the integration between the host and the virtual machine:

- **Graceful Shutdown**—Allows you to perform a graceful shutdown of the VM-Series firewall from the Hyper-V management interface without having to log into the guest.
- **Heartbeat to Hyper-V Manager**—Provides heartbeat monitoring of the running status of guest VMs from the Hyper-V management interface.
- **Firewall Management IP Address Visibility**—Allows you to use Hyper-V Manager to view the IP address assigned to the management interface on the firewall.

 ©2024 Palo Alto Networks, Inc.

Install the VM-Series Firewall on Hyper-V

Use the instructions in this section to deploy your VM-Series firewall on a Hyper-V host. A Palo Alto Networks support account and a valid VM-Series license are required to download the VHDX image file and install the VM-Series on the Hyper-V host. If you have not already registered the capacity auth-code that you received with the order fulfillment email, with your support account, see Register the VM-Series Firewall. After completing the registration continue to the following tasks:

- Before You Begin
- Performance Tuning of the VM-Series Firewall on Hyper-V
- Provision the VM-Series Firewall on a Hyper-V host with Hyper-V Manager
- Provision the VM-Series Firewall on a Hyper-V host with PowerShell
- Perform Initial Configuration on the VM-Series Firewall

Before You Begin

Before installing and configuring your VM-Series firewall, know and account for the following items as needed when you configure your VM-Series firewall:

- Virtual Switch Types
- MAC Address Spoofing

Virtual Switch Types

Before installing the VM-Series, you must create the vSwitches required for providing external connectivity for management access and for routing traffic from and to the virtual machines that the firewall will secure. Hyper-V allows you to create three types of vSwitches:

- **External vSwitch**—binds to a physical network adapter and provides the vSwitch access to a physical network.
- **Internal vSwitch**—passes traffic between the virtual machines and the Hyper-V host. This type of vSwitch does not provide connectivity to a physical network connection.
- **Private vSwitch**—passes traffic between the virtual machines on the Hyper-V host only.

An external vSwitch is required for management of the VM-Series firewall. Other vSwitches connected to the VM-Series firewall can be of any type and will depend on your network topology.

MAC Address Spoofing

If you are deploying the VM-Series firewall with interfaces enabled in Layer 3 mode, make sure to enable use of hypervisor assigned MAC addresses so that the hypervisor and the firewall can properly handle packets. Alternatively, use the Hyper-V Manager to enable MAC address spoofing on the virtual network adapter for each dataplane interface on the firewall. For more information, see Hypervisor Assigned MAC Addresses.

If you are deploying the VM-Series firewall with interfaces enabled in Layer 2 mode or virtual-wire mode, you must enable MAC address spoofing on the virtual network adapter in Hyper-V for each dataplane interface on the firewall. This setting is required to ensure that packets sent by the

 ©2024 Palo Alto Networks, Inc.

VM-Series are not dropped by the virtual network adapter if the source MAC address does not match the outgoing interface MAC address.

Performance Tuning of the VM-Series Firewall on Hyper-V

The VM-Series firewall for Hyper-V is a high-performance appliance but may require tuning of the hypervisor to achieve the best results. This section describes some best practices and recommendations for facilitating the best performance of the VM-Series firewall.

- Disable Virtual Machine Queues
- Isolate CPU Resources in a NUMA Node

Disable Virtual Machine Queues

Palo Alto Networks recommends disabling virtual machine queues (VMQ) for all NICs on the Hyper-V host. This option is prone to misconfiguration and can cause reduced network performance when enabled.

STEP 1 | Login to Hyper-V Manager and select your VM.

STEP 2 | Select **Settings** > **Hardware** > **Network Adapter** > **Hardware Acceleration**.

STEP 3 | Under Virtual machine queue, uncheck **Enable virtual machine queue**.

STEP 4 | Click **Apply** save your changes and **OK** to exit the VM settings.

Isolate CPU Resources in a NUMA Node

You can improve performance of VM-Series for Hyper-V by isolating the CPU resources of the guest VM to a single non-uniform memory access (NUMA) node. You can view the NUMA settings of your VM in Hyper-V Manager by selecting **Settings** > **Hardware** > **Processor** > **NUMA**.

Provision the VM-Series Firewall on a Hyper-V host with Hyper-V Manager

Use these instructions to deploy the VM-Series firewall on Hyper-V using Hyper-V Manager.

STEP 1 | Download the VHDX file.

Register your VM-Series firewall and obtain the VHDX file.

1. Go to https://www.paloaltonetworks.com/services/support.
2. Filter by **PAN-OS for VM-Series Base Images** and download the VHDX file. For example, PA-VM-HPV-7.1.0.vhdx.

STEP 2 | Set up any vSwitch(es) that you will need.

To create a vSwitch:

1. From Hyper-V Manager, select the host and select **Action** > **Virtual Switch Manager** to open the Virtual Switch Manager window.
2. Under **Create virtual switch**, select the type of vSwitch (external, internal, or private) to create and click **Create Virtual Switch**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Install the firewall.

1. On the Hyper-V Manager, select the host and select **Action** > **New** > **Virtual Machine.** Configure the following settings in the New Virtual Machine Wizard:
   1. Choose a **Name** and **Location** for the VM-Series firewall. The VM-Series firewall stores the VHDX file at the specified location.
   2. Choose **Generation 1**. This is the default option and the only version supported.
   3. For **Startup Memory**, assign the memory based on the VM-Series System Requirements of your VM-Series model.

      Do not enable dynamic memory; the VM-Series firewall requires static memory allocation.

   4. Configure **Networking**. Select an external vSwitch to connect the management interface on the firewall.
   5. To connect the **Virtual Hard Disk**, select **Use an existing virtual hard disk** and browse to the VHDX file you downloaded earlier.
   6. Review the summary and click **Finish**.
2. Assign virtual CPUs to the firewall.
   1. Select the VM you created and navigate to **Action** > **Settings**.
   2. Select **Processor** and enter the minimum number of CPUs based on the VM-Series System Requirements of your VM-Series model..
   3. Click **OK**.

STEP 4 | Connect at least one network adapter for the dataplane interface on the firewall.

1. Select **Settings** > **Hardware** > **Add Hardware** and select the **Hardware type** for your network adapter.

   Legacy Network Adapter and SR-IOV are not supported. If selected, the VM-Series firewall will boot into maintenance mode.

2. Click **OK**.

STEP 5 | (Optional) Enable MAC address spoofing on Hyper-V if you are not using Layer 3 with hypervisor assigned MAC address.

1. Double click the dataplane virtual network adapter and click **Advanced Settings**.
2. Click the **Enable MAC address spoofing** check box and click **Apply**.

STEP 6 | Power on the firewall.

Select the firewall from the list of **Virtual Machines** and navigate to **Action** > **Start** to power on the firewall.

Provision the VM-Series Firewall on a Hyper-V host with PowerShell

Use these instructions to deploy the VM-Series firewall on Hyper-V using PowerShell.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Download the VHDX file.

Register your VM-Series firewall and obtain the VHDX file.

1. Go to https://www.paloaltonetworks.com/services/support.
2. Filter by **PAN-OS for VM-Series Base Images** and download the VHDX file. For example, PA-VM-HPV-7.1.0.vhdx.

STEP 2 | Set up any vSwitch(es) that you will need.

Create a vSwitch by using the following commands. Give the vSwitch a name and choose the switch type.

```
> New-VMSwitch -Name <"switch-name"> -SwitchType <switch-type>
```

STEP 3 | Install the VM-Series firewall.

1. Create the new virtual machine and set the memory based on the VM-Series System Requirements of your VM-Series model.

   ```
   > NEW-VM -Name <vm-name> -MemoryStartupBytes 4GB -
   VHDPath <file-path-to-vhdx>
   ```

2. Set processor count based on the VM-Series System Requirements of your VM-Series model.

   ```
   > SET-VMProcessor –VMName <vm-name> –Count 2
   ```

STEP 4 | Connect at least one network adapter for the management interface on the firewall.

Connect the default network adapter created during VM creation to management vSwitch.

```
> connect-VMNetworkAdapter -vmname <vm-name> -Name <"network-
adapter-name"> -SwitchName <"management-vswitch">
```

STEP 5 | (Optional) Enable MAC address spoofing on Hyper-V if you are not using Layer 3 with hypervisor assigned MAC address.

```
> Set-VMNetworkAdapter -vmname <vm-name> -Name <"network-adapter-
name"> -MacAddressSpoofing On
```

STEP 6 | Power on the firewall.

For example:

```
> Start-VM -vmname <vm-name>
```

 ©2024 Palo Alto Networks, Inc.

Perform Initial Configuration on the VM-Series Firewall

Use these instructions to perform the initial configuration of your VM-Series firewall. By default, the VM-Series firewall uses DHCP to obtain an IP address for the management interface. However, you can assign a static IP address. After completing the initial configuration, access the web interface to complete further configurations tasks. If you have Panorama for central management, refer to the Panorama Administrator's Guide for information on managing the device using Panorama.

If you are using bootstrapping to perform the configuration of your VM-Series firewall on Hyper-V, refer to Bootstrap the VM-Series Firewall on Hyper-V. For general information about bootstrapping, see Bootstrap the VM-Series Firewall.

STEP 1 | Gather the required information from your network administrator.

- Management port IP address
- Netmask
- Default gateway
- DNS server IP address

STEP 2 | Access the console of the VM-Series firewall.

1. In Hyper-V Manager, select the VM-Series firewall and click **Connect** from the Actions list.
2. Log in to the firewall with the default username and password: **admin/admin**
3. Enter configuration mode using the following command: **configure**

STEP 3 | Configure the network access settings for the management interface.

Enter the following commands:

```
set deviceconfig system type static
```

```
set deviceconfig system ip-address <Firewall-IP> netmask <netmask>
default-gateway <gateway-IP> dns-settingservers primary <DNS-IP>
```

where *<Firewall-IP>* is the IP address you want to assign to the management interface, *<netmask>* is the subnet mask, *<gateway-IP>* is the IP address of the network gateway, and *<DNS-IP>* is the IP address of the DNS server.

STEP 4 | Commit your changes and exit the configuration mode.

1. Enter **commit**.
2. Enter **exit**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Verify that you can view the management interface IP address from the Hyper-V Manager.

1. Select the VM-Series firewall from the list of **Virtual Machines**.
2. Select **Networking**. The first network adapter that displays in the list is used for management access to the firewall; subsequent adapters in the list are used as the dataplane interfaces on the firewall.

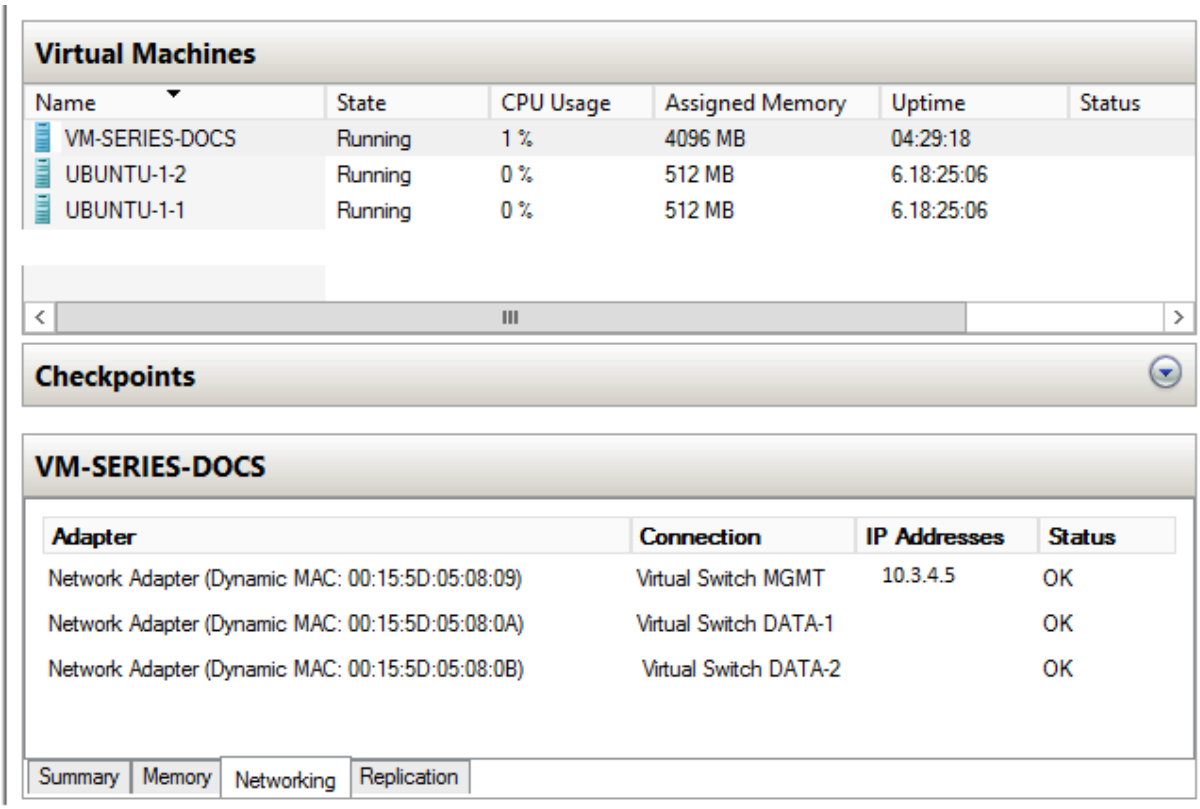

STEP 6 | Verify network access to external services required for firewall management, such as the Palo Alto Networks Update Server.

1. Use the ping utility to verify network connectivity to the Palo Alto Networks Update server as shown in the following example. Verify that DNS resolution occurs and the response includes the IP address for the Update server; the update server does not respond to a ping request.

```
admin@PA-220 > ping host updates.paloaltonetworks.com
```

```
PING updates.paloaltonetworks.com (10.101.16.13) 56(84) bytes
 of data.
From 192.168.1.1 icmp_seq=1 Destination Host Unreachable
From 192.168.1.1 icmp_seq=2 Destination Host Unreachable
From 192.168.1.1 icmp_seq=3 Destination Host Unreachable
```

 ©2024 Palo Alto Networks, Inc.

```
From 192.168.1.1 icmp_seq=4 Destination Host Unreachable
```

After verifying DNS resolution, press Ctrl+C to stop the ping request.

2. Use the following CLI command to retrieve information on the support entitlement for the firewall from the Palo Alto Networks update server:

```
request support
check
```

If you have connectivity, the update server will respond with the support status for your firewall.

STEP 7 | (Optional) Verify that your VM-Series jumbo frame configuration does not exceed the maximum MTU supported on Hyper-V.

The VM-Series has a default MTU size of 9216 bytes when jumbo frames are enabled. However, the maximum MTU size supported by the physical network adapter on the Hyper-V host is 9000 or 9014 bytes depending on the network adapter capabilities. To verify the configured MTU on Hyper-V:

1. In Windows Server 2012 R2, open the **Control Panel** and navigate to N**etwork and Internet** > **Network and Sharing Center** > **View network status and tasks**.
2. Click on a network adapter or virtual switch from the list.
3. Click **Properties**.
4. Click **Configure**.
5. On the Advanced tab, select **Jumbo Packet** from the list.
6. Select 9000 or 9014 bytes from the Value drop-down menu.
7. Click **OK**.

If you have enabled jumbo frames on Hyper-V, Enable Jumbo Frames on the VM-Series Firewall and set the MTU size to match that configured on the Hyper-V host.

STEP 8 | Access the web interface of the VM-Series firewall and configure the interfaces and define security rules and NAT rules to safely enable the applications you want to secure.

Refer to the PAN-OS Administrator's Guide.

 ©2024 Palo Alto Networks, Inc.

Set up the VM-Series Firewall on Azure

VM-Series firewall on Azure brings the security features of Palo Alto Networks next generation firewall as a virtual machine in the Azure Marketplace. The VM-Series firewall provides a complete set of security functionality to ensure that your virtual machine workloads and data are protected, and the capabilities that the firewall enables are different from native security features such as Security Groups, Web Application Firewalls and native, port-based firewalls.

On Azure, the VM-Series firewall is available in the bring your own license (BYOL) model or in the pay-as-you-go (PAYG) hourly model. Microsoft Azure allows you to deploy the firewall to secure your workloads within the virtual network in the cloud, so that you can deploy a public cloud solution or you can extend the on-premises IT infrastructure to create a hybrid solution.

- About the VM-Series Firewall on Azure
- Deployments Supported on Azure
- Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template)
- Deploy the VM-Series Firewall from the Azure China Marketplace (Solution Template)
- Panorama Orchestrated Deployments in Azure
- Create a Custom VM-Series Image for Azure
- Use Azure Security Center Recommendations to Secure Your Workloads
- Deploy the VM-Series Firewall on Azure Stack
- Enable Azure Application Insights on the VM-Series Firewall
- VM Monitoring on Azure
- Set up Active/Passive HA on Azure
- Use the ARM Template to Deploy the VM-Series Firewall
- Deploy the VM-Series and Azure Application Gateway Template
- Secure Kubernetes Services on Azure

About the VM-Series Firewall on Azure

The VM-Series firewall on Azure must be deployed in a virtual network (VNet) using the Resource Manager deployment mode. You can deploy the VM-Series firewall on the standard Azure public cloud, Azure China, and Azure Government—including DoD on Azure Government, which meets the security requirements for DoD Impact Level 5 data and FedRAMP High standards.

The VM-Series firewall on the marketplace for the Azure public cloud, Azure Government, and Azure DoD regions, supports both the Bring Your Own License (BYOL) model and the hourly Pay-As-You-Go (PAYG) option (usage-based licensing) For licensing details, see License Types—VM-Series Firewalls, and refer to the list of supported Azure regions in which you can deploy the VM-Series firewall.

For Azure China, the VM-Series firewall is available in the BYOL option only. See Deploy the VM-Series Firewall from the Azure China Marketplace (Solution Template) for the workflow.

You can also deploy the VM-Series firewall on Azure Stack, Microsoft's private cloud solution that enables you to use Azure services within your organization's datacenter. With Azure Stack, you can build out a hybrid cloud solution that unifies your public Azure deployment with your on-premise Azure Stack set up. You can download the VM-Series firewall BYOL offer from the Azure Marketplace and make it available to your tenants on Azure Stack. For instructions, see Deploy the VM-Series Firewall on Azure Stack.

- Azure Networking and VM-Series Firewall
- Azure Security Center Integration
- VM-Series Firewall Templates on Azure
- Minimum System Requirements for the VM-Series on Azure
- Support for High Availability on VM-Series on Azure

Azure Networking and VM-Series Firewall

The Azure VNet infrastructure does not require virtual machines to have a network interface in each subnet. The architecture includes an internal route table (called system routes) that directly connects all virtual machines within a VNet such that traffic is automatically forwarded to a virtual machine in any subnet. For a destination IP address that is not within the VNet, the traffic is sent to the default Internet gateway or to a VPN gateway, if configured. In order to route traffic through the VM-Series firewall, you must create user defined routes (UDRs) that specify the next hop for traffic leaving a subnet. This route forces traffic destined to another subnet to go to the VM-Series firewall instead of using the system routes to directly access the virtual machine in the other subnet. For example, in a two-tiered application with a web tier and a database tier, you can set up UDRs for directing traffic from the web subnet to the DB subnet through the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

On Azure, UDRs are for traffic leaving a subnet only. You cannot create user defined routes to specify how traffic comes into a subnet from the Internet or to route traffic to virtual machines within a subnet. UDRs allow you to direct outbound traffic to an interface on the VM-Series firewall so that you can always ensure that the firewall secures traffic to the internet also.

For documentation on Microsoft Azure, refer to https://azure.microsoft.com/en-us/documentation/.

The solution templates for deploying the VM-Series firewall that are available in the Azure Marketplace, have three network interfaces. To Set up Active/Passive HA on Azure, you will need to add an additional interface for the HA2 link. If you want to customize the template, use the ARM templates that are available in the GitHub repository.

Azure Security Center Integration

Microsoft has deprecated Azure Security Center support for partner security solutions and replaced it with Azure Sentinel.

The VM-Series firewall is integrated with Azure Security Center to provide a unified view for monitoring and alerting on the security posture of your Azure workloads. On Azure Security Center, the VM-Series firewall is available as a partner security solution that protects your Azure workloads from threats and mitigates any gaps in securing your business and intellectual property in the public cloud. To enable this integration and display logs as security alerts directly on the Azure Security Center dashboard, the VM-Series firewall on Azure includes a Log Forwarding profile.

To get started, you need to enable Azure Security Center on your Azure subscription. You then have two ways in which you can enable this integration:

- **Deploy the VM-Series firewall based on a recommendation on the Azure Security Center dashboard.**

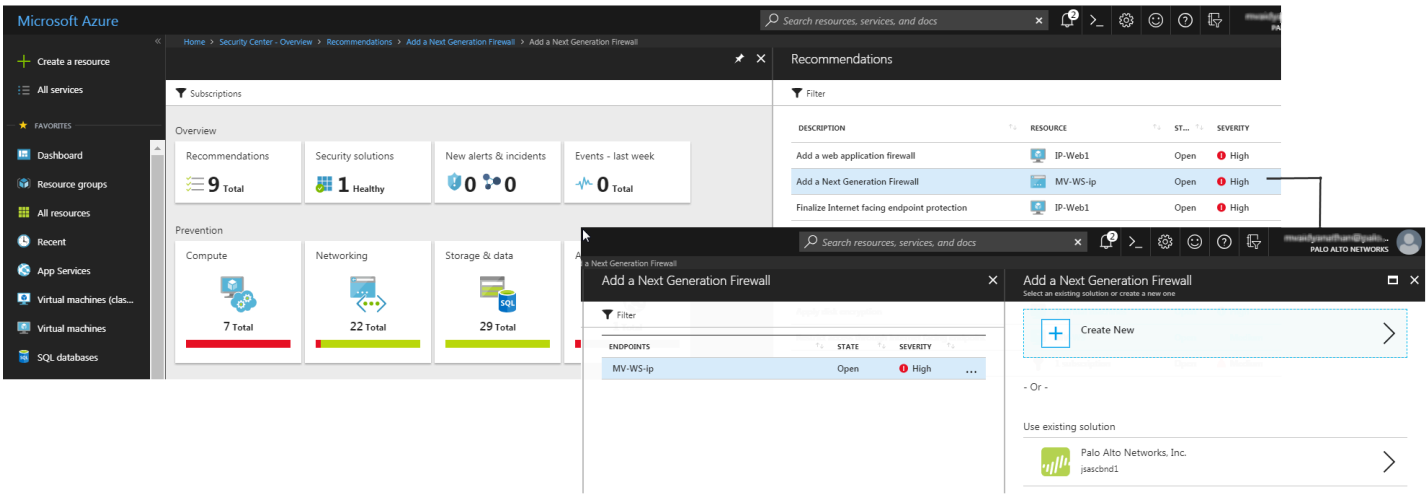

When the Azure Security Center dashboard recommends that you deploy a VM-Series firewall to secure a workload that is exposed to the internet, you can only deploy the firewall in an new resource group or an existing resource group that is empty. This is because Azure currently restricts you from deploying a multi NIC appliance in an existing resource group. Therefore,

 ©2024 Palo Alto Networks, Inc.

after you deploy the VM-Series firewall you must manually configure it to be in the path of traffic of the workload that you need to secure.

When you deploy the firewall from Azure Security Center, the firewall is launched with three network interfaces—management, external facing (untrust) and internal facing (trust)—and a user defined route (UDR) that sends all outbound traffic from the trust subnet to the trust interface on the firewall so that internet-bound traffic is always inspected by the firewall. The default configuration includes two example Security Policy rules—the *outbound-default* rule allows all traffic from the trust zone to the untrust zone on the application default port, and the *inbound-default* rule allows all web-browsing traffic from the untrust zone to the trust zone, after inspecting traffic with the default Antivirus, Anti-spyware, and Vulnerability Protection security profiles. The firewall also forwards all files that are intercepted with the inbound or outbound rule to the WildFire public cloud for analysis. Both rules include a URL Filtering profile that blocks all traffic to the URL categories copyright-infringement, dynamic-dns, extremism, malware, phishing, and unknown. In addition to these security profiles, both Security policy rules are enabled to log at session end and to forward Threat and WildFire Submissions logs as security alerts to the Azure Security Center dashboard.

To make practical use of this integration and Deploy a VM-Series Firewall Based on an Azure Security Center Recommendationwithin the same resource group as the workloads you want to secure, you can stage a workload with a public IP address that is exposed to the internet. When Azure Security Center detects the security risk, it triggers a recommendation to deploy a next-generation firewall, and you can then deploy the VM-Series firewall in a new resource group into which you can add your workloads later. You must then delete the workload that you staged to trigger the recommendation.

- **Select a VM-Series firewall that you have already deployed for securing your workloads**. If you have a Standard tier of Azure Security Center subscription, Azure Security Center discovers and displays all existing VM-Series firewalls that you have deployed either from the Azure Marketplace or using a customized deployment with Azure CLI, PowerShell or ARM

 ©2024 Palo Alto Networks, Inc.

template. The firewalls within your Azure subscription are grouped under Security Solutions on the Azure Security Center dashboard.

Microsoft Azure does not support the discovery of existing firewalls with the Free tier subscription.

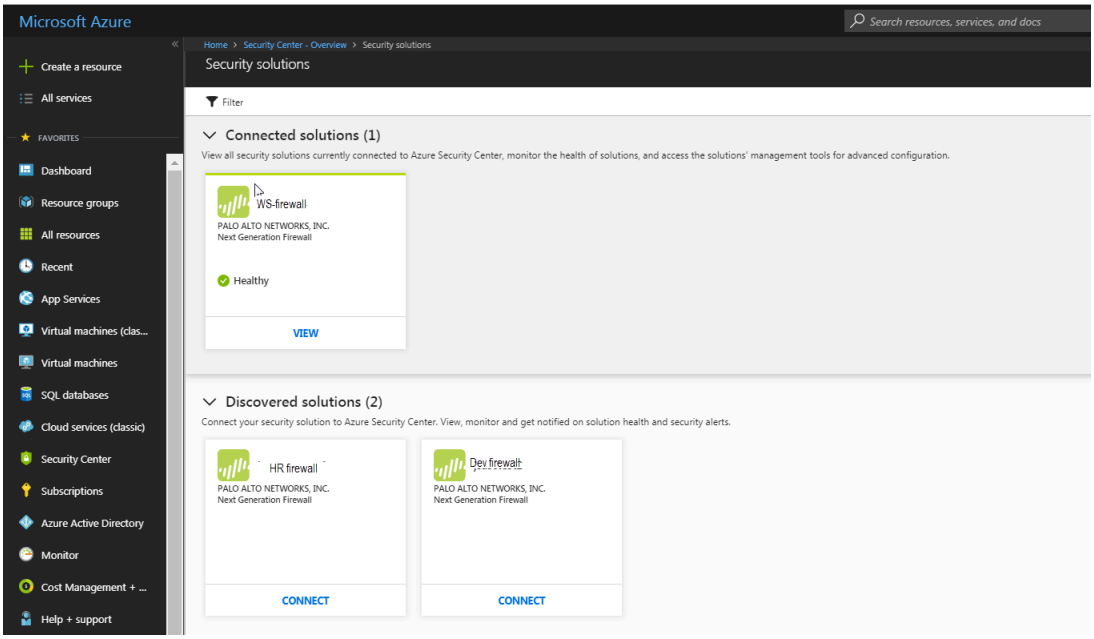

To Connect an Existing VM-Series Firewall From Azure Security Center, you must set up a Linux virtual machine and configure Syslog forwarding to forward firewall logs in the Common Event Format as alerts to Azure Security Center. The additional configuration enables a single pane of glass view for monitoring all your Azure assets.

Forwarding a large volume of logs to Azure Security Center, may result in additional subscription cost to you.

VM-Series Firewall Templates on Azure

You can deploy the VM-Series firewall on Azure using templates. Palo Alto Networks provides two kinds of templates—Solution templates and ARM templates.

- **Solution Templates in the Azure Marketplace** —The solution templates that are available in the Azure Marketplace allow you to deploy the VM-Series firewall using the Azure portal. You can use an existing resource group and storage account (or create them new) to deploy the VM-Series firewall with the following default settings for all regions except Azure China:
  - VNet CIDR 10.8.0.0/16; you can customize the CIDR to a different private IP address range.
  - Three subnets— 10.8.0.0/24 (management), 10.8.1.0/24 (untrust), 10.8.2.0/24 (trust)
  - Three network interfaces, one in each subnet. If you customize the VNet CIDR, the subnet ranges map to your changes.

  To use the solution template, see Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template) for Azure China, see Deploy the VM-Series Firewall from the Azure China Marketplace (Solution Template).

 ©2024 Palo Alto Networks, Inc.

- **ARM Templates in the GitHub Repository**—In addition to Marketplace based deployments, Palo Alto Networks provides Azure Resource Manager templates in the GitHub Repository to simplify the process of deploying the VM-Series firewall on Azure.
  - Use the ARM Template to Deploy the VM-Series Firewall—The basic ARM template includes two JSON files (a Template file and a Parameters File) to help you deploy and provision all the resources within the VNet in a single, coordinated operation. These templates are provided under an as-is, best effort, support policy.

    If you want to use the Azure CLI to locate all the images available from Palo Alto Networks, you the need the following details to complete the command (show vm-image list):

    - *Publisher: paloaltonetworks*
    - *Offer: vmseries-flex*
    - *SKU: byol, bundle1, bundle 2*
    - *Version: 10.0.0, or latest*
  - Deploy the VM-Series and Azure Application Gateway Template to support a scale out security architecture that protects your internet-facing web applications using two VM-Series firewalls between a pair of (external and internal) Azure load balancers VM-Series and Azure Application Gateway. This template is currently not available for Azure China.
  - Use the ARM template to deploy the VM-Series firewall in to an existing Resource Group, for example when you want to Set up Active/Passive HA on Azure.

In addition to the ARM templates above that are covered under the Palo Alto Networks official support policy, Palo Alto Networks provides Community supported templates in the Palo Alto Networks GitHub repository that allow you to explore the solutions available to jumpstart your journey in to cloud automation and scale on Azure.

Minimum System Requirements for the VM-Series on Azure

You must deploy the VM-Series firewall in the Azure Resource Manager (ARM) mode only; the classic mode (Service Management based deployments) is not supported. The VM-Series firewall on Azure must meet the following requirements:

- Azure Linux VMs of the following types—Supported models.

  These types include support for Accelerated Networking (SR-IOV).
- For memory, disk and CPU cores required to deploy the VM-Series firewall, see VM-Series System Requirements.

  You can add additional disk space of 40GB to 8TB for logging purposes. The VM-Series firewall uses Azure managed disks where available; it does not utilize the temporary disk that Azure provides with some instance types.
- Up to eight network interfaces (NICs). A primary interface is required for management access and up to seven interfaces for data traffic.

  On Azure, because a virtual machine does not require a network interface in each subnet, you can set up the VM-Series firewall with three network interfaces (one for management traffic and two for dataplane traffic). To create zone-based policy rules on the firewall, in addition to the management interface, you need at least two dataplane interfaces so that you can assign

 ©2024 Palo Alto Networks, Inc.

one dataplane interface to the *trust* zone, and the other dataplane interface to the *untrust* zone. For an HA deployment, you will need another interface for the HA2 link between the HA peers.

Because the Azure VNet is a Layer 3 network, the VM-Series firewall on Azure supports Layer 3 interfaces only.

Support for High Availability on VM-Series on Azure

To ensure availability, you can Set up Active/Passive HA on Azurein a traditional configuration with session synchronization, or use a scale out architecture using cloud-native load balancers such as the Azure Application Gateway or Azure Load Balancer to distribute traffic across a set of healthy instances of the firewall. For details, see Deploy the VM-Series and Azure Application Gateway Template.

VM-Series on Azure Service Principal Permissions

For Panorama to interact with the Azure APIs and collect information on your workloads, you need to create an Azure Active Directory application and a Service Principal that has the permissions required to authenticate with Azure AD and access the resources within your subscription.

To create the Active Directory application and Service Principal, follow the instructions in How to: Use the portal to create an Azure AD application and service principal that can access resources. During the application generation process, there is a step to "Assign application to role" and assign an IAM role of "reader" to the application.

If you don't have the necessary permissions to create and register the AD application, ask your Azure AD or subscription administrator to create a Service Principal.

After the application has been registered, record these values so you can enter them in the Panorama plugin for Azure at a later time:

- Application ID
- Secret Key (record it when you make the secret key; the secret key is not visible once you navigate away from the page).
- Tenant ID

Permissions

The following table lists the minimum built-in roles required and the granular permissions if you would like to customize the role.

To support	Permissions
Azure High Availability	See Set up Active/Passive HA on Azure.
Azure Application Insights	**“Microsoft.Authorization/*/read”,** **“Microsoft.Network/networkInterfaces/*”,** **“Microsoft.Network/networkSecurityGroups/*”,**

 ©2024 Palo Alto Networks, Inc.

To support	Permissions
Enable Azure Application Insights on the VM-Series Firewall	**"Microsoft.Network/virtualNetworks/*",** **"Microsoft.Compute/virtualMachines/read"**
Azure Monitoring Set Up the Azure Plugin for Monitoring on Panorama	Requires a minimum Role of **Reader** for Service Principal. Alternatively, you can add the following custom permissions: **"Microsoft.Compute/virtualMachines/read",** **"Microsoft.Network/networkInterfaces/read",** **"Microsoft.Network/virtualNetworks/read",** **"Microsoft.Network/locations/serviceTags/read"** **"Microsoft.Network/loadBalancers/read",** **"Microsoft.Network/publicIPAddresses/read"** **"Microsoft.Resources/subscriptions/resourcegroups/read",**
Panorama Orchestrated Deployments Create a Custom Role and Associate it with an Active Directory	**"Microsoft.Resources/subscriptions/resourcegroups/*",** **"Microsoft.Resources/deployments/write",** **"Microsoft.Resources/deployments/operationStatuses/read",** **"Microsoft.Resources/deployments/read",** **"Microsoft.Resources/deployments/delete"**
	"Microsoft.Network/publicIPPrefixes/write", **"Microsoft.Network/publicIPPrefixes/read",** **"Microsoft.Network/publicIPPrefixes/delete",** **"Microsoft.Network/publicIPAddresses/write",** **"Microsoft.Network/publicIPAddresses/read",** **"Microsoft.Network/publicIPAddresses/delete",** **"Microsoft.Network/publicIPAddresses/join/action",**
	"Microsoft.Network/natGateways/write", **"Microsoft.Network/natGateways/read",** **"Microsoft.Network/natGateways/delete",** **"Microsoft.Network/natGateways/join/action",**

 ©2024 Palo Alto Networks, Inc.

To support	Permissions
	"Microsoft.Network/virtualNetworks/read", "Microsoft.Network/virtualNetworks/write", "Microsoft.Network/virtualNetworks/delete", "Microsoft.Network/virtualNetworks/subnets/write", "Microsoft.Network/virtualNetworks/subnets/read", "Microsoft.Network/virtualNetworks/subnets/delete", "Microsoft.Network/virtualNetworks/subnets/join/action",
	"Microsoft.Network/virtualNetworks/virtualNetworkPeerings/read",
	"Microsoft.Network/networkSecurityGroups/write", "Microsoft.Network/networkSecurityGroups/read", "Microsoft.Network/networkSecurityGroups/delete", "Microsoft.Network/networkSecurityGroups/join/action",
	"Microsoft.Network/loadBalancers/write", "Microsoft.Network/loadBalancers/read", "Microsoft.Network/loadBalancers/delete", "Microsoft.Network/loadBalancers/probes/join/action", "Microsoft.Network/loadBalancers/backendAddressPools/join/action", "Microsoft.Network/loadBalancers/frontendIPConfigurations/read",
	"Microsoft.Network/locations/serviceTags/read",
	"Microsoft.Network/applicationGateways/read", "Microsoft.Network/networkInterfaces/read",
	"Microsoft.Compute/virtualMachineScaleSets/write",

 ©2024 Palo Alto Networks, Inc.

To support	Permissions
	"Microsoft.Compute/virtualMachineScaleSets/read", "Microsoft.Compute/virtualMachineScaleSets/delete", "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/read",
	"Microsoft.Compute/virtualMachines/read",
	"Microsoft.Compute/images/read",
	"Microsoft.insights/components/write", "Microsoft.insights/components/read", "Microsoft.insights/components/delete",
	"Microsoft.insights/autoscalesettings/write"

 ©2024 Palo Alto Networks, Inc.

Deployments Supported on Azure

Use the VM-Series firewall on Azure to secure your network users in the following scenarios:

- **Hybrid and VNet to VNet**—The VM-Series firewall on Azure allows you to securely extend your physical data center/private cloud into Azure using IPSec and ExpressRoute. To improve your data center security, if you have segmented your network and deployed your workloads in separate VNets, you can secure traffic flowing between VNets with an IPSec tunnel and policies that allow application traffic.

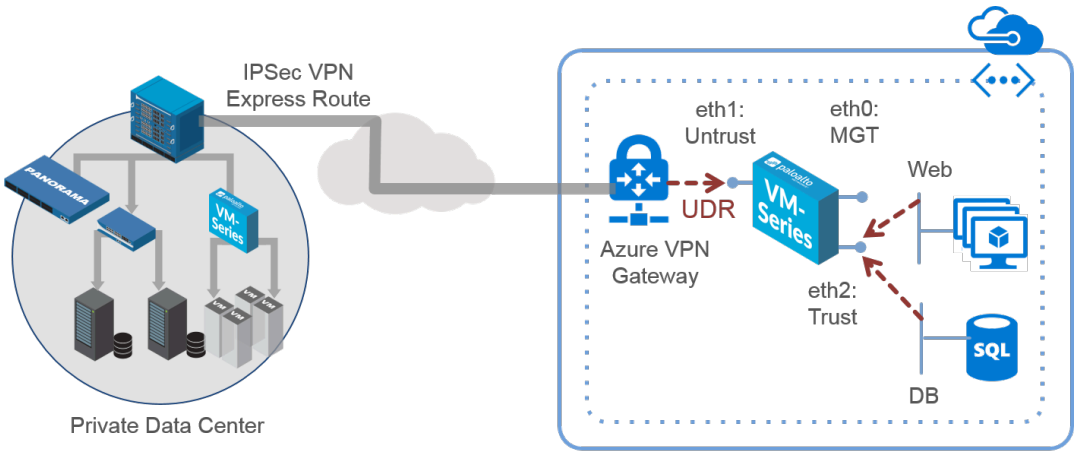

- **Inter-Subnet** —The VM-Series firewall can front your servers in a VNet and protects against lateral threats for inter-subnet traffic between applications in a multi-tier architecture.
- **Gateway**—The VM-Series firewall serves as the VNet gateway to protect Internet-facing deployments in the Azure Virtual Network (VNet). The VM-Series firewall secures traffic destined to the servers in the VNet and it also protects against lateral threats for inter-subnet traffic between applications in a multi-tier architecture.
- **GlobalProtect**—Use the Azure infrastructure to quickly and easily deploy the VM-Series firewall as GlobalProtect™ and extend your gateway security policy to remote users and devices, regardless of location.

You can continue with Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template), Deploy the VM-Series Firewall on Azure Stack, or Orchestrate a VM-Series Firewall Deployment in Azure.

You can also learn about the VM-Series Firewall Templates on Azure that you can use to deploy the firewall.

For information on bootstrapping, see Bootstrap the VM-Series Firewall on Azure.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template)

The following instructions describe how to deploy the solution template for the VM-Series firewall that is available in the Azure® Marketplace and the Azure Government Marketplace. To use the customizable Azure Resource Manager (ARM) templates available in the GitHub repository, see Use the ARM Template to Deploy the VM-Series Firewall.

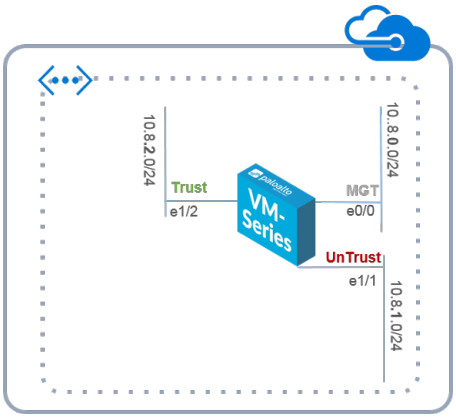

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Set up an Azure account.

1. If you don't have one already, create a Microsoft® account.
2. Log in to the Azure portal (**`https://portal.azure.com`** or **`https://portal.azure.us`**) using your Microsoft account credentials.

   *If you are using a trial subscription, you may need to open a support request (**Help + Support** > **New Support Request**) to increase the quota of allocated VM cores.*

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Find the VM-Series solution template in the Azure Marketplace.

1. Select **Marketplace** > **Virtual Machines**.
2. Search for Palo Alto Networks® and a list of offerings for the VM-Series firewall will display. For the differences in the BYOL (bring your own license) and PAYG (pay as you go) models, see VM-Series Firewall Licenses for Public Clouds.

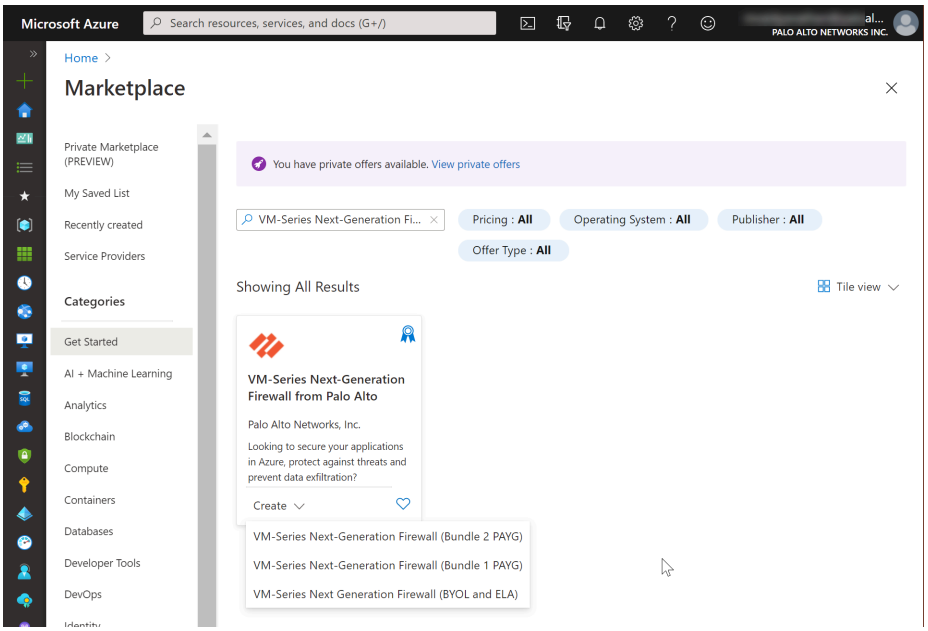

3. Select an offering to **Create** a new VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Deploy the firewall.

1. Configure basic settings for the firewall.
   1. Select your Azure **Subscription**.
   2. Create a new resource group or select an existing resource group that is empty. The resource group will hold all the resources associated with the VM-Series firewall for this deployment.

      Azure has removed the option to select an existing resource group for Marketplace solutions that enable multiple network interface controllers (NICs). To deploy the firewall into an existing resource group, use the ARM template in the GitHub Repository or use your own custom ARM template.
   3. Select the Azure **Region** in which you are deploying the firewall.
   4. Enter a **Username** for the firewall administrator.
   5. Select the **Authentication type**—Password or SSH Public Key.

      You must enable SSH key authentication if you plan to use the firewall in FIPS-CC operational mode. Although you can deploy the VM-Series firewall using a username and password, you will be unable to authenticate using the username and password after changing the operational mode to FIPS-CC. After resetting to FIPS-CC mode, you must use the SSH key to log in and can then configure a username and password that you can use for subsequently logging in to the firewall web interface. For details on creating the SSH key, refer to the Azure documentation.
   6. Enter a **Password** (up to 31 characters) or copy and paste an **SSH public key** for securing administrative access to the firewall.
2. Configure networking.
   1. Select an existing Azure Virtual Network (VNet) or create a new one and enter the IP address space for the VNet. By default, the Classless Inter-Domain Routing (CIDR) IP address is 10.8.0.0/16.
   2. Configure the subnets for the network interfaces.

      If you use the default subnets, you must review the configuration. If you use an existing VNet, you must have set up three subnets: one each for the management,trust, and untrust interfaces. If you create a new VNet, verify or change the prefixes for each subnet. The default subnets are 10.8.0.0/24 for the

 ©2024 Palo Alto Networks, Inc.

management subnet, 10.8.1.0/24 for the untrust subnet, and 10.8.2.0/24 for the trust subnet.

3. Enter the source IP address or IP range (include the CIDR block) that can access the VNet. **Network Security Group: inbound source IP** allows you to restrict inbound access to the Azure VNet.

   Restrict access to the firewall. Make sure to supply a CIDR block that corresponds to your dedicated management IP addresses or network. Do not make the allowed source network range larger than necessary and never configure the allowed source as 0.0.0.0/0. Verify your IP address before you configure it on the template to make sure that you do not lock yourself out.

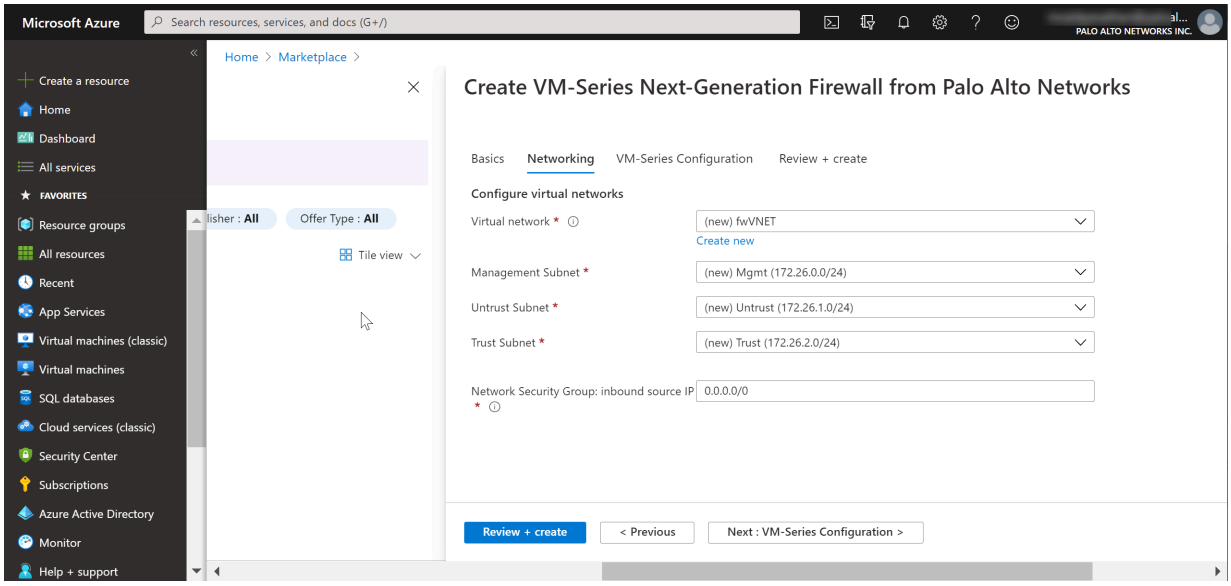

3. Define management access to the firewall.

   1. Use the default variable ((new) fwMgmtPublicIP)) to assign a **Public IP address** to the management interface (`eth0`) of the firewall.

      Azure accelerated networking is not supported on the management interface.

   2. Enter a prefix to access the firewall using a DNS name. You must combine the prefix you enter with the suffix displayed on screen to access the web interface of the firewall. For example: **`<yourname><your-region>.cloudapp.azure.com`**
   3. Select latest **VM-Series Version**.
   4. Enter a display name to identify the VM-Series firewall within the resource group.

 ©2024 Palo Alto Networks, Inc.

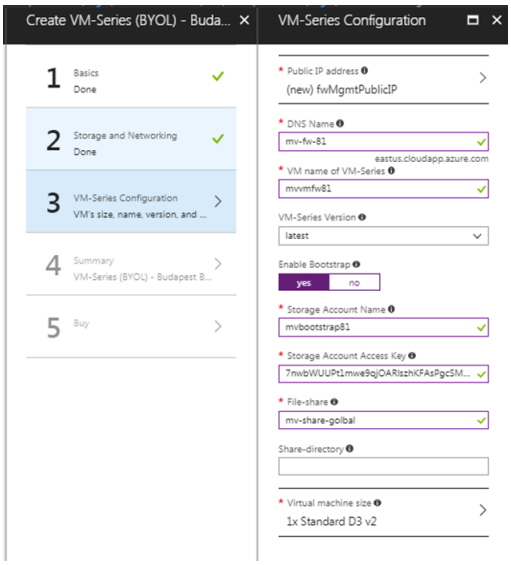

4. Add the information to configure the firewall at launch. See Bootstrap the VM-Series Firewall on Azure.
   1. Select **yes** to **Enable Bootstrap**.
   2. Enter the **Storage Account Name** that holds the Bootstrap Package.
   3. Enter the **Storage Account Access Key**. This firewall needs this access key to authenticate to the storage account and access the files stored within.
   4. Add the **File share name** to which you have uploaded the files required for bootstrapping the firewall. The storage account must be in the same region in which

 ©2024 Palo Alto Networks, Inc.

you are deploying the firewall and it must have the correct folder structure for bootstrapping.

5. Select the Azure virtual machine tier and size to meet your needs. Use the **Change size** link to view supported instance types, and to review the Minimum System Requirements for the VM-Series on Azure.

5. Review the summary, and **OK**. Then accept the terms of use and privacy policy, and **Create** to launch the firewall.

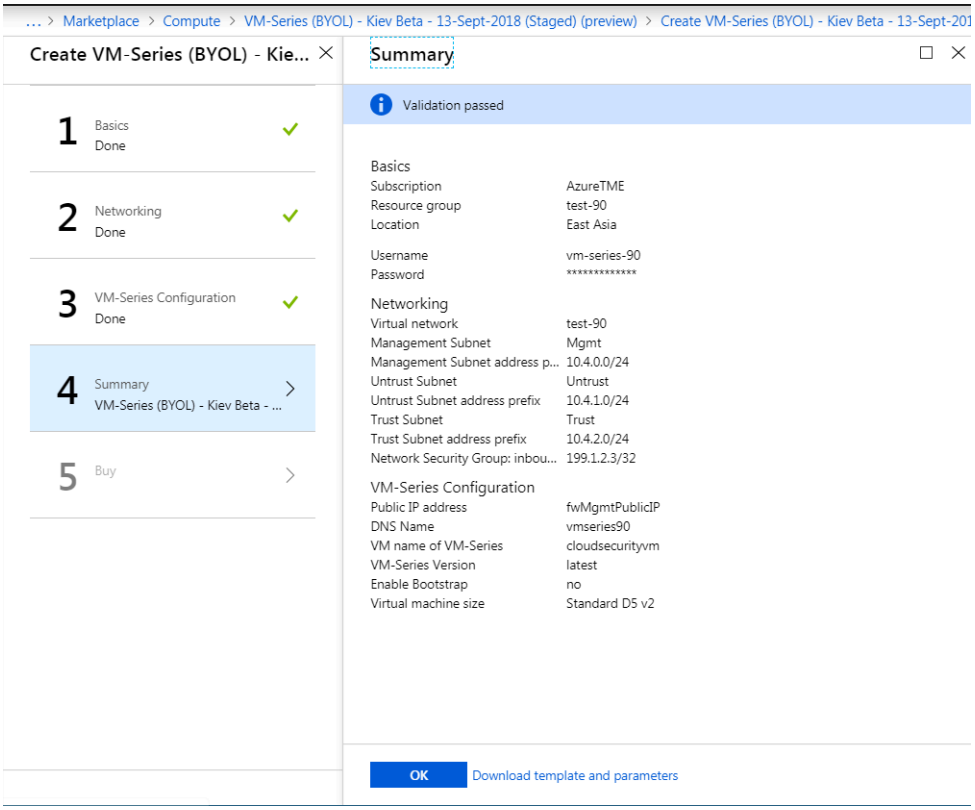

6. Verify that you have successfully deployed the VM-Series firewall.

   1. Select **Dashboard** > **Resource Groups** and select the resource group.
   2. Select your resource group and see the **Overview** for detailed status on which resources are deployed successfully.

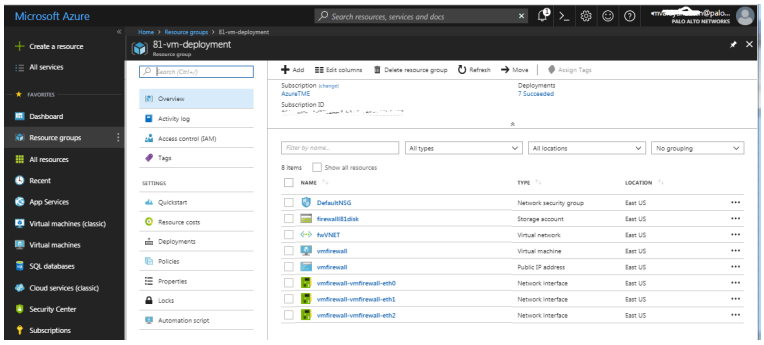

STEP 4 | Attach a public IP address for the untrust interface of the VM-Series firewall. When you create a new public IP address, you get one from the block of IP addresses that Microsoft

 ©2024 Palo Alto Networks, Inc.

owns, so you can't choose a specific one. The maximum number of public IP addresses you can assign to an interface is based on your Azure subscription.

1. On the Azure portal, select the network interface for which you want to add a public IP address (such as the **eth1** interface).
2. Select **IP Configurations** > **Add** and, for Public IP address, select **Enabled**. Create a new public IP address or select one that you have available.
3. Verify that you can view the secondary IP address associated with the interface.

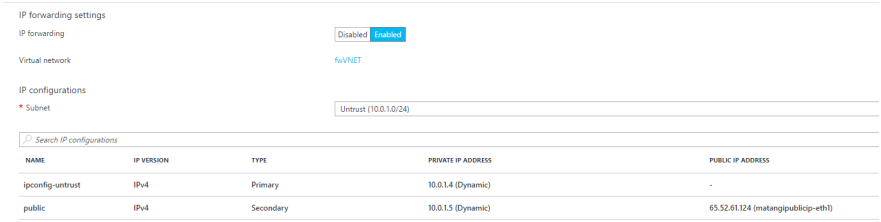

When you attach a secondary IP address to a network interface, the VM-Series firewall does not automatically acquire the private IP address assigned to the interface. You will need to manually configure the private IP address using the VM-Series firewall web interface. See Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

STEP 5 | Log in to the web interface of the firewall.

1. On the Azure portal, in **All Resources**, select the VM-Series firewall and view the full DNS name for the firewall.

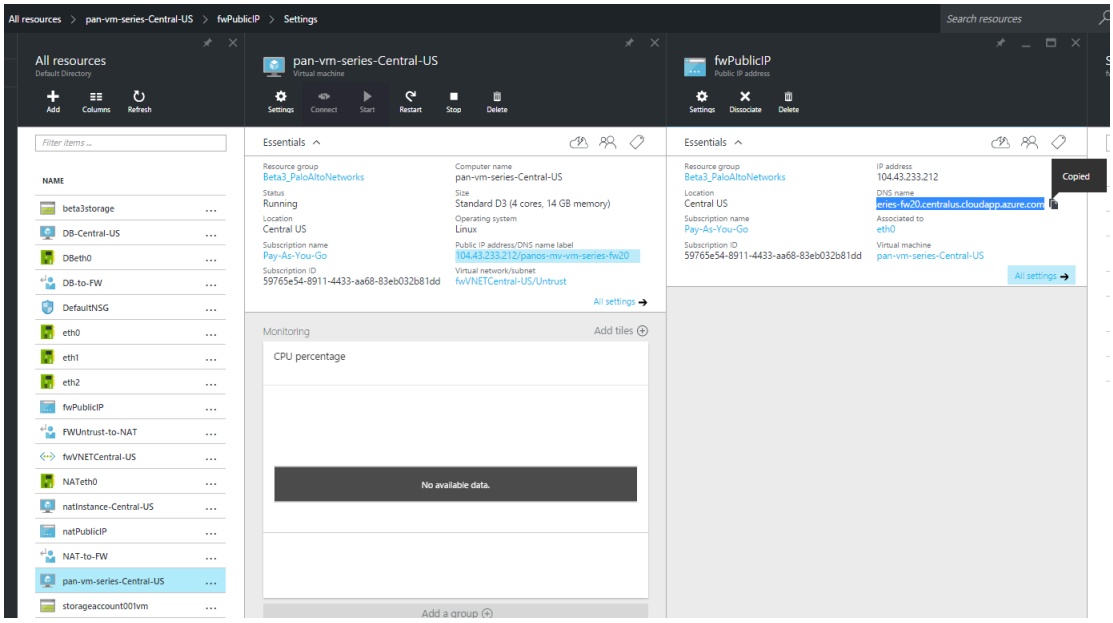

1. Using a secure (https) connection from your web browser, log in to the DNS name for the firewall.
2. Enter the usernamepassword that you defined in the parameters file. You will see a certificate warning but that is OK—continue to the web page.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Activate the licenses on the VM-Series firewall.

For the BYOL version

1. Create a Support Account.
2. Register the VM-Series Firewall (with auth code).
3. On the firewall web interface, select **Device** > **Licenses** and select **Activate feature using authentication code**.
4. Enter the capacity authentication code (*auth-code*) that you registered on the support portal. The firewall will connect to the update server (`updates.paloaltonetworks.com`), and download the license and reboot automatically.
5. Log back in to the web interface and confirm the following on the **Dashboard**:
   - A valid serial number displays in **Serial#**.

     If the term `Unknown` displays, it means the firewall is not licensed. To view traffic logs on the firewall, you must install a valid capacity license.
   - The **VM Mode** displays as Microsoft Azure.

For the PAYG version

1. Create a Support Account.
2. Register the Usage-Based Model of the VM-Series Firewall for Public Clouds (no auth code).

STEP 7 | Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

If you are hosting multiple websites or services with different IP addresses and SSL certificates on a single server, you might need to configure more than one IP address on the VM-Series firewall interfaces.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Click **ethernet 1/1** and configure as follows:
   - Set **Interface Type** to **Layer3** (default).
   - On the **Config** tab, assign the interface to the default router.
   - Also on the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone called **`UnTrust`**, and then click **OK**.
   - On the **IPv4** tab, select **DHCP Client** if you plan to assign only one IP address on the interface—the firewall will automatically acquire the private IP address assigned in the ARM template. If you plan to assign more than one IP address, select **Static** and

 ©2024 Palo Alto Networks, Inc.

manually enter the primary and secondary IP addresses assigned to the interface on the Azure portal.

- Disable (clear) the **Automatically create default route to default gateway provided by server** to ensure that traffic handled by this interface does not flow directly to the default gateway in the VNet.

3. Click **ethernet 1/2** and configure as follows:
   - Set **Interface Type** to **Layer3** (default).
   - Set **Security Zone** to **Trust**.
   - Set **IP address DHCP Client** or **Static**.
   - Disable (clear) the **Automatically create default route to default gateway provided by server**to ensure that traffic handled by this interface does not flow directly to the default gateway in the VNet.
4. **Commit** your changes and verify that the link state for the interfaces is up.
5. Add a static route on the virtual router of the VM-Series firewall for any networks that the firewall needs to route.

   For example, to add a default route to the destination subnets for the servers that the firewall secures:
   - Select **Network** > **Virtual Router** > **default** >
   - Select **Static Routes** > **IPv4**, and add the next hop IP address for the destination servers. You can set x.x.x.1 as the next hop IP address for all traffic (destined to 0.0.0.0/0 from interface ethernet1/1).

STEP 8 | Configure the firewall for your specific deployment.

- **Gateway**—Deploy a third-party load balancer in front of the UnTrust zone.
- **Hybrid and Inter-VNet**—Deploy an Azure VPN Gateway or a NAT virtual machine in front the UnTrust zone.
- **Inter-Subnet**—On the VM-Series firewall, add an intrazone Security policy rule to allow traffic based on the subnets attached to the Trust interface.
- **GlobalProtect™**—Deploy a NAT virtual machine in front of the UnTrust zone.

STEP 9 | Direct traffic to the VM-Series firewall.

1. To ensure that the VM-Series firewall secures all traffic within the Azure resource group, configure static routes on the firewall.
2. Configure user defined routes to direct all traffic through the interfaces on the VM-Series firewall. Refer to the Azure documentation on UDRs for details.

   The user defined routes on the internal subnets must send all traffic through the Trust interface. The user defined routes on the UnTrust side direct all traffic from the Internet through the UnTrust interface on the VM-Series firewall. The traffic from the Internet may be coming from an Azure Application Gateway or Azure Load Balancer, or through the Azure VPN Gateway in the case of a hybrid deployment that connects your on-premise network with the Azure cloud.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | To publish PAN-OS® metrics to Azure Application Insights, see Enable Azure Application Insights on the VM-Series Firewall.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall from the Azure China Marketplace (Solution Template)

The following instructions show you how to deploy the solution template for the VM-Series firewall that is available in the Azure China Marketplace. The Azure China Marketplace supports only the BYOL model of the VM-Series firewall. You can deploy the firewall in a existing resource group that is empty or into a new resource group. The default VNet in the template is 10.0.0.0/16, and it deploys a VM-Series firewall has 3 network interfaces, one management and two dataplane interfaces as shown below. To use the customizable ARM templates available in the GitHub repository, see Use the ARM Template to Deploy the VM-Series Firewall.

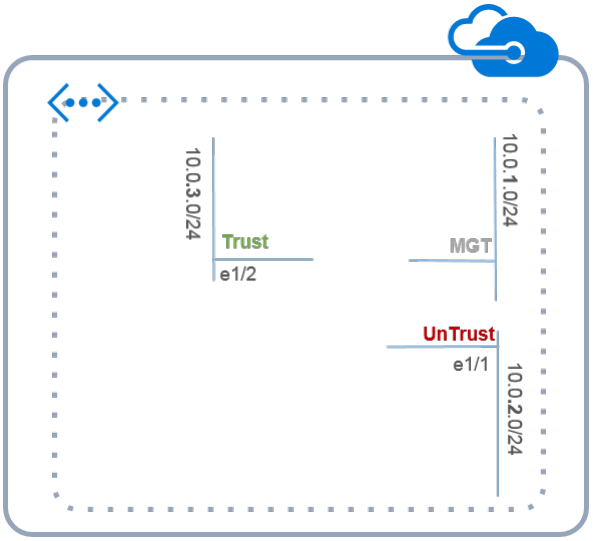

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Set up an Azure account.

1. Create a Microsoft account.
2. Log in to the Azure portal (**`https://portal.azure.com`**) using your Microsoft account credentials.

   *If you are using a trial subscription, you may need to open a support request (**Help + Support** > **New Support Request**) to increase the quota of allocated VM cores.*

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Find the VM-Series solution template in the Azure Marketplace.

1. Search for Palo Alto Networks on the Azure China marketplace (https://market.azure.cn). The offering for the different PAN-OS versions of the VM-Series firewalls displays.

2. Select an offering and click **Immediate deployment of**.

STEP 3 | Deploy the firewall.

1. Select your Azure **Subscription**.
2. Select a resource group for holding all the resources associated with the VM-Series firewall in this deployment.

   You can deploy the VM-Series firewall into a new Resource Group, or an existing Resource Group that is empty. To deploy the firewall into an existing resource group that has other resources, use the ARM template in the GitHub Repository or your own custom ARM template. Ensure that the existing resources match the parameter values you provide in the ARM template.

   1. If you create a new resource group, enter a name for the resource group and select the Azure China region where you want to deploy the firewall.
   2. If you select an existing resource group, select the Azure China region for this resource group, and select complete deployment.
3. Configure basic settings for the firewall.
   1. Enter the storage account name for an existing account or create a new one.
   2. Enter the name for the blob storage container to which the firewall vhd mage will be copied and saved.
   3. Enter a DNS name for accessing the Public IP address on the management interface (eth0) of the firewall. To access the web interface of the firewall, you must combine the prefix you enter with the suffix, for example <yourDNSname><china_region>.cloudapp.azure.com.
   4. Enter a **Username** for the firewall administrator.
   5. Enter a **Password** for securing administrative access to the firewall.
   6. Select the Azure virtual machine tier and size to meet your needs. See Minimum System Requirements for the VM-Series on Azure.
   7. Enter a **VmName**, which is a display name to identify the VM-Series firewall within the resource group.
   8. Use a **PublicIPAddressName** to label the firewall management interface within the resource group. Microsoft Azure binds the DNS name that you defined with this

 ©2024 Palo Alto Networks, Inc.

name so that you can access the management interface on the firewall from the public internet.

9. Enter a **VirtualNetworkName** to identify your VNet. The default IP **Address Prefix** for the VNet is 10.0.0.0/16. You can change this to meet your IP addressing needs.
10. Configure the subnets for the network interfaces. If you use an existing VNet, you must have defined three subnets, one each for the management, trust and untrust interfaces. If you create a new VNet, verify or change the prefixes for each subnet. The default subnets are 10.0.1.0/24, 10.0.2.0/24, and 10.0.3.0/24. You can allocate these subnets to the management, trust, and untrust interfaces as you would like.

4. Review the summary, accept the terms of use and privacy policy, and click **Immediate deployment** to deploy the firewall. The deployment maybe take 20 minutes and you can use the link on the page to verify progress.
5. Verify that you have successfully deployed the VM-Series firewall.
   1. Log in to the Azure China portal (https://portal.azure.cn) using your Microsoft account credentials.
   2. Select **Dashboard** > **Resource Groups**, select the resource group.
   3. Select **All Settings** > **Deployments** > **Deployment History** for detailed status

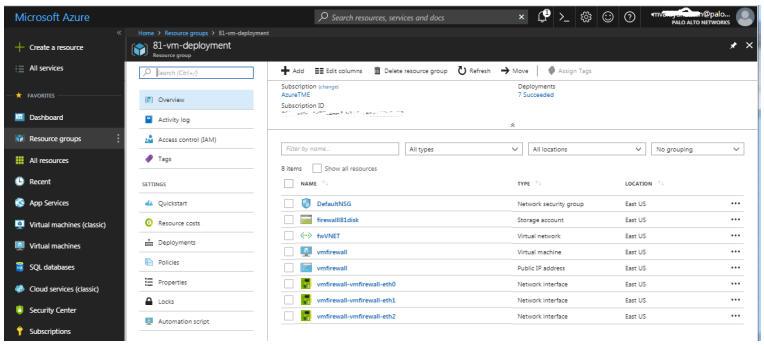

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Attach a public IP address for the untrust interface of the VM-Series firewall. This allows you to access the interface from the public internet and is useful for any internet-facing application or service.

1. On the Azure portal, select the network interface for which you want to add a public IP address. For example the eth1 interface.
2. Select **IP Configurations** > **Add** and for Public IP address, select **Enabled**. Create a new public IP address or select one that you have available.
3. Verify that you can view the secondary IP address associated with the interface.

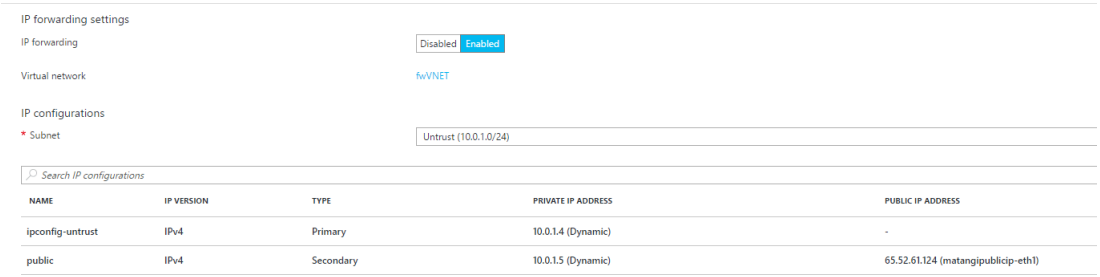

When you attach a secondary IP address to a network interface, the VM-Series firewall does not automatically acquire the private IP address assigned to the interface. You will need to manually configure the private IP address using the VM-Series firewall web interface. See Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

Each interface on the VM-Series firewall on Azure can have one dynamic (default) or static private IP address, and multiple public IP addresses (static or dynamic) associated with it. The maximum number of public IP addresses you can assign to an interface is based on your Azure subscription. When you create a new public IP address you get one from the block of IP addresses Microsoft owns, so you can't choose a specific one.

STEP 5 | Log in to the web interface of the firewall.

1. On the Azure portal, in **All Resources**, select the VM-Series firewall and view the full DNS name for the firewall.

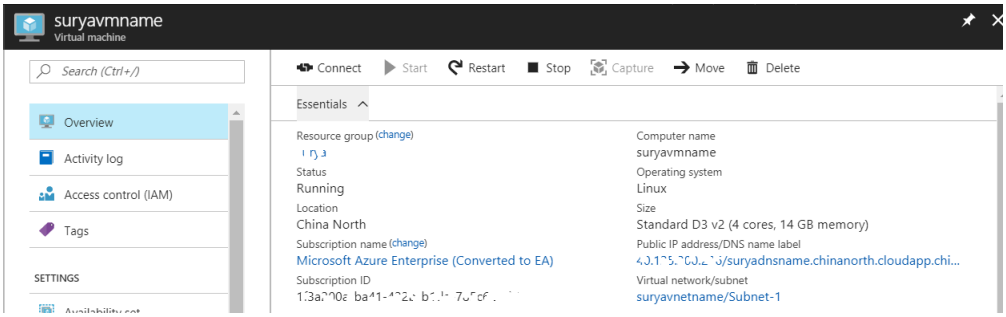

2. Using a secure connection (https) from your web browser, log in to the DNS name for the firewall.
3. Enter the username/password you defined earlier. You will see a certificate warning; that is okay. Continue to the web page.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Activate the licenses on the VM-Series firewall.

1. Create a Support Account.
2. Register the VM-Series Firewall (with auth code).
3. On the firewall web interface, select **Device** > **Licenses** and select **Activate feature using authentication code**.
4. Enter the capacity auth-code that you registered on the support portal. The firewall will connect to the update server (updates.paloaltonetworks.com), and download the license and reboot automatically.
5. Log back in to the web interface and confirm the following on the **Dashboard**:
   - A valid serial number displays in **Serial#**.

     If the term Unknown displays, it means the device is not licensed. To view traffic logs on the firewall, you must install a valid capacity license.
   - The **VM Mode** displays as Microsoft Azure.

STEP 7 | Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

If you are hosting multiple websites or services with different IP addresses and SSL certificates on a single server, you might need to configure more than one IP address on the VM-Series firewall interfaces.

1. Select **Network** > **Interfaces** > **Ethernet**.
2. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: Layer3 (default).
   - On the Config tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select New Zone. Define a new zone called UnTrust, and then click OK.
   - On the **IPv4** tab, select **DHCP Client** if you plan to assign only one IP address on the interface. The private IP address assigned in the ARM template will be automatically acquired. If you plan to assign more than one IP address select **Static** and manually enter the primary and secondary IP addresses assigned to the interface on the Azure portal.
   - Clear the **Automatically create default route to default gateway provided by server** check box. Disabling this option ensures that traffic handled by this interface does not flow directly to the default gateway in the VNet.
3. Click the link for **ethernet 1/2** and configure as follows:
   - Set **Interface Type** to Layer3 (default).
   - **Security Zone**: Trust
   - **IP address**: Select **DHCP Client** or **Static**.
   - Clear the **Automatically create default route to default gateway provided by server** check box. Disabling this option ensures that traffic handled by this interface does not flow directly to the default gateway in the VNet.
4. Click **Commit**. Verify that the link state for the interfaces is up.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Configure the firewall for your specific deployment.

- Gateway—Deploy a 3rd party load balancer in front of the UnTrust zone.
- Hybrid and Inter-VNet—Deploy an Azure VPN Gateway or a NAT virtual machine in front the UnTrust zone.
- Inter-Subnet—On the VM-Series firewall, add an intra-zone security policy rule to allow traffic based on the subnets attached to the Trust interface.
- GlobalProtect—Deploy a NAT virtual machine in front of the UnTrust zone.

STEP 9 | Direct traffic to the VM-Series firewall.

1. To ensure that the VM-Series firewall secures all traffic within the Azure resource group, configure static routes on the firewall.
2. Configure UDRs to direct all traffic through the interfaces on the VM-Series firewall. Refer to the Azure documentation on UDRs for details.

   The UDRs on the internal subnets must send all traffic through the Trust interface. The UDRs on the UnTrust side direct all traffic from the Internet through the UnTrust interface on the VM-Series firewall. The traffic from the Internet may be coming from an Azure Application Gateway or Azure Load Balancer, or through the Azure VPN Gateway in case of a hybrid deployment that connects your on-premises network with the Azure cloud.

 ©2024 Palo Alto Networks, Inc.

Panorama Orchestrated Deployments in Azure

The Panorama plugin for Azure centrally deploys, configures, and monitors your security posture in Azure cloud. It orchestrates VM-Series deployments in your Azure network so that you can enable security policies for managed firewalls. The plugin links to your Azure ARM deployment and Azure Monitor pages, providing visibility into the deployment status, usage, and performance of your VM-Series firewalls.

In Azure, the plugin orchestrates the deployment of Azure resources such as load balancers, subnets and NAT gateways as well as VM-Series firewall autoscaling sets. In Panorama the plugin automatically configures Panorama device groups, template stacks, and NAT policies. It reads the tags from your Azure resources, then centrally enables tag-based policies on a group of firewalls.

The Panorama plugin can orchestrate deployments in one or more regions in your Azure environment. A deployment can consist of a hub stack or an inbound stack or both, depending on the traffic that needs to be secured for your deployment:

- A *Hub firewall* stack protects outbound traffic and East-West traffic between your application workloads.
- An *Inbound firewall* stack secures traffic to and from your public facing applications.

You can configure the number of firewalls in each stack. You have the option to configure a static amount of firewalls in your deployment or a range for the VMSS to use for scaling. Both stacks in the deployment create a VMSS of VM-Series firewalls and they can each scale up to as many as 25 firewalls.

Hub Stack

A deployment uses a Hub stack and leverages the Azure Internal Standard Load Balancer (with HA ports) to scale and load balance across a set of firewalls. You can then use the Standard Load balancer's private IP address (**2**, "Hub/Egress Private IP" in the following figure) to route traffic to the firewalls for inspection and threat prevention. The Hub stack secures your applications' outbound and East-West traffic.

 ©2024 Palo Alto Networks, Inc.

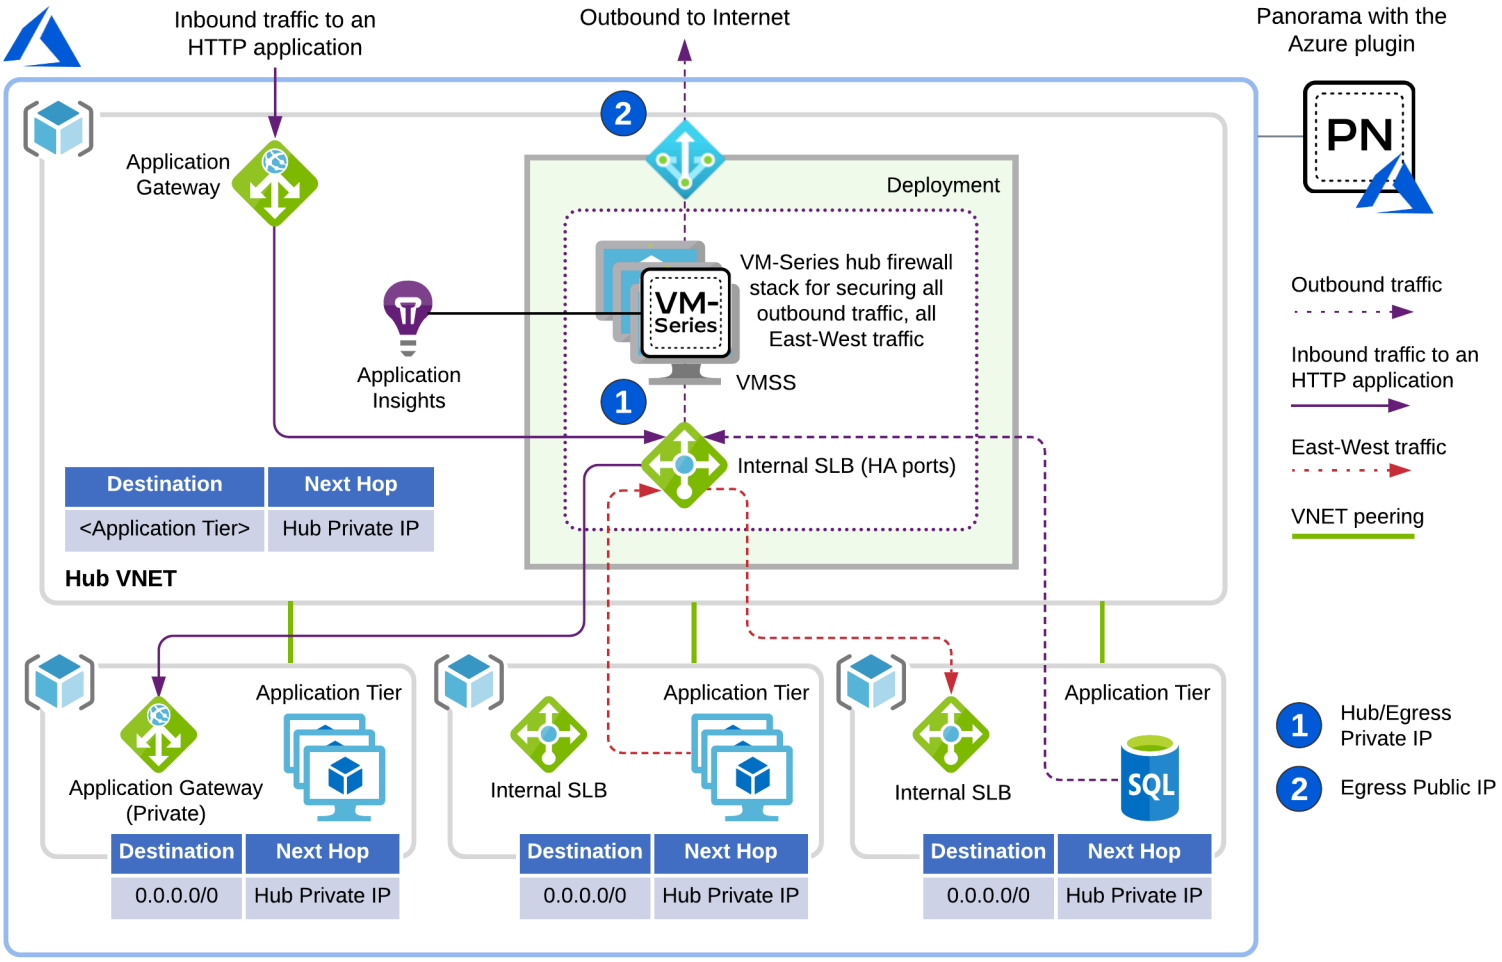

To protect your outbound traffic and East-West traffic, add route rules in your application VNETs to redirect traffic to the Hub stack for inspection.

Inbound Stack

An Inbound firewall stack scales independently and adds visibility and security to your applications' Inbound traffic.

Each inbound stack can secure up to 10 applications.

To protect your inbound HTTP traffic, add UDRs in the Application Gateway's subnet route tables to route all traffic to the Inbound stack (**3**, Ingress Private IP in the following figure). To protect the non-HTTP inbound traffic, use the Panorama plugin to create front-end entries for your application endpoints (**4**, Ingress Public IP Front Ends in the following figure). To enable inspection, the Panorama plugin automatically creates load balancer rules on the Azure Public Standard Load Balancer and NAT rules on the firewalls.

If you only have HTTP/HTTPS inbound traffic you can leave out the Inbound stack and protect that traffic with just the hub stack.

 ©2024 Palo Alto Networks, Inc.

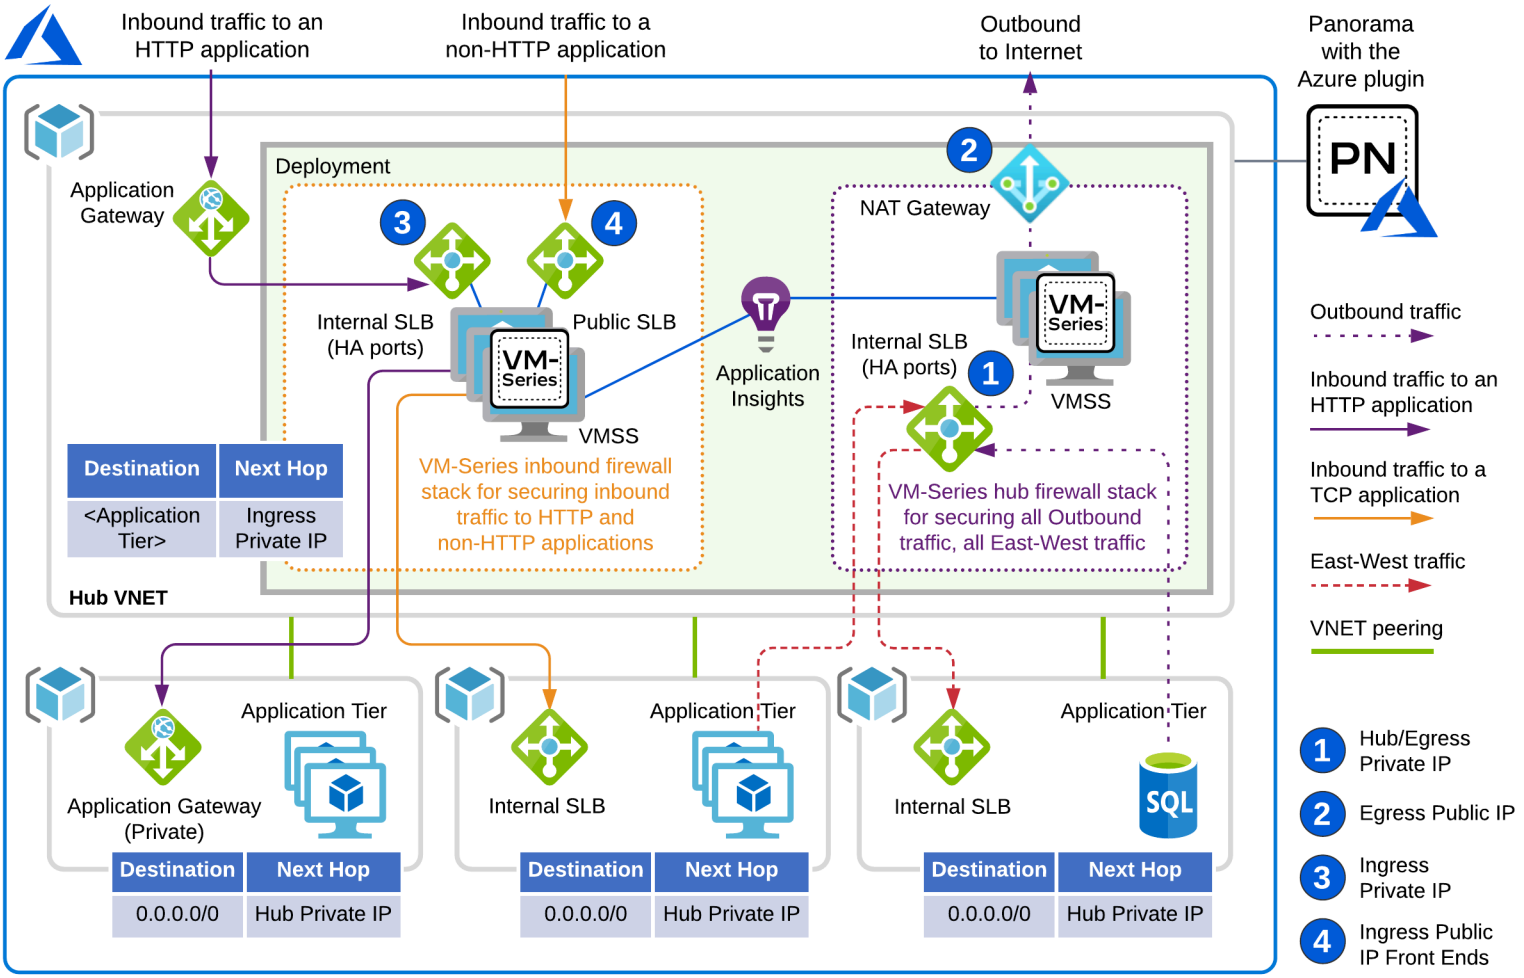

See Prepare for an Orchestrated Deployment and Orchestrate a VM-Series Firewall Deployment in Azure.

Prepare for an Orchestrated Deployment

Complete the following tasks before you orchestrate a VM-Series firewall on Azure.

- Configuration Prerequisites
- Orchestration Permissions
- Create a Custom Role and Associate it with an Active Directory
- Find Your Azure Directory Domain Name

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

Configuration Prerequisites

Complete the following basic tasks on Panorama and Azure.

 ©2024 Palo Alto Networks, Inc.

- Azure
  - Create a service principal to enable the plugin to make API calls.
  - Plan a CIDR block specifically dedicated to VM-Series firewall Transit VNet. The plugin manages this CIDR block and uses it to deploy the initial firewall VNet and perform as future upgrades to new template stacks.

    The minimum CIDR range is /22.
- Panorama
  - Ensure that you have a valid license API key configured on Panorama. This allows the plugin to manage delicensing on scale-in autoscaling events. See Install a License Deactivation API Key.
  - Install the latest version of the VM-series plugin on Panorama to allow Application Insight configuration to be added to template stack.
- Qualified Azure Regions

  Panorama orchestrated deployments are supported in all regions that support the VM-Series firewall. The following regions have been qualified; if you deploy in an unlisted region and you encounter an issue, contact support.
  - West US
  - West US 2
  - North Central US
  - East US
  - East US 2
  - West Europe
  - Germany West Central
  - UAE North
  - West India
  - Australia Southeast

While planning your deployment please note that if you are currently running a Panorama Plugin for Azure version 2.x, upgrading to the current version is not allowed. Along with this, once the current version is installed, downgrading to a version 2.x plugin is not allowed. See Panorama Plugin for Azure in the Compatibility Matrix.

Orchestration Permissions

The is a sample JSON file with permissions for the Template Deployer role. In the **`AssignableScopes`** section, include all relevant subscriptions that must be queried, including the subscription into which the deployment is deployed and EVERY subscription containing an application VNET that is peered to the VM-Series firewall VNet where protected resources exist.

```
{
    "Name": "Template Deployment",
    "IsCustom": true,
    "Description": "Manage template deployments.",
    "Actions": [
```

 ©2024 Palo Alto Networks, Inc.

```
        "Microsoft.Resources/subscriptions/resourcegroups/*",
        "Microsoft.Resources/deployments/write",
        "Microsoft.Resources/deployments/operationStatuses/read",
        "Microsoft.Resources/deployments/read",
        "Microsoft.Resources/deployments/delete",
        "Microsoft.Network/publicIPPrefixes/write",
        "Microsoft.Network/publicIPPrefixes/read",
        "Microsoft.Network/publicIPPrefixes/delete",
        "Microsoft.Network/publicIPAddresses/write",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Network/publicIPAddresses/delete",
        "Microsoft.Network/publicIPAddresses/join/action",
        "Microsoft.Network/natGateways/write",
        "Microsoft.Network/natGateways/read",
        "Microsoft.Network/natGateways/delete",
        "Microsoft.Network/natGateways/join/action",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/virtualNetworks/write",
        "Microsoft.Network/virtualNetworks/delete",
        "Microsoft.Network/virtualNetworks/subnets/write",
        "Microsoft.Network/virtualNetworks/subnets/read",
        "Microsoft.Network/virtualNetworks/subnets/delete",
        "Microsoft.Network/virtualNetworks/subnets/join/action",
        "Microsoft.Network/virtualNetworks/virtualNetworkPeerings/
read",
        "Microsoft.Network/networkSecurityGroups/write",
        "Microsoft.Network/networkSecurityGroups/read",
        "Microsoft.Network/networkSecurityGroups/delete",
        "Microsoft.Network/networkSecurityGroups/join/action",
        "Microsoft.Network/loadBalancers/write",
        "Microsoft.Network/loadBalancers/read",
        "Microsoft.Network/loadBalancers/delete",
        "Microsoft.Network/loadBalancers/probes/join/action",
        "Microsoft.Network/loadBalancers/backendAddressPools/join/
action",
        "Microsoft.Network/loadBalancers/frontendIPConfigurations/
read",
        "Microsoft.Network/locations/serviceTags/read",
        "Microsoft.Network/applicationGateways/read",
        "Microsoft.Network/networkInterfaces/read",
        "Microsoft.Compute/virtualMachineScaleSets/write",
        "Microsoft.Compute/virtualMachineScaleSets/read",
        "Microsoft.Compute/virtualMachineScaleSets/delete",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/
read",
        "Microsoft.Compute/virtualMachines/read",
        "Microsoft.Compute/images/read",
        "Microsoft.Compute/virtualMachines/read",
        "Microsoft.Network/networkInterfaces/read",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/locations/serviceTags/read",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/loadBalancers/read",
        "Microsoft.insights/components/write",
        "Microsoft.insights/components/read",
        "Microsoft.insights/components/delete",
```

 ©2024 Palo Alto Networks, Inc.

```
            "Microsoft.insights/autoscalesettings/write"
        ]
         "NotActions": [
        ],
         "AssignableScopes": [
            "/subscriptions/{deployment-subscription}",
            "/subscriptions/{app1-subscription}",
            "/subscriptions/{app2-subscription}",
            .
            .
            .
        ]
}
```

Create a Custom Role and Associate it with an Active Directory

STEP 1 | To create an Active Directory in Azure, Navigate to Azure Active Directory and click on **App Registrations** on the left, then click **Add** and provide a name.

STEP 2 | Create a custom role with the permissions the plugin requires.

See Orchestration Permissions.

1. Log in to the Azure CLI.

   ```
   az login
   ```

2. Create a custom role from the file in Orchestration Permissions.

   ```
   az role definition create --role-definition <role-json-file>
   ```

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Associate the role with Active Directory you created in Step 1. You can use the console or the CLI.

You must repeat this step in every subscription defined in the ***assignableScope*** *section of the custom role in Orchestration Permissions.*

Console

1. On the Azure portal, navigate to **Subscriptions** and select your subscription.
2. On the left select **Access Control (IAM)** and then **Role Assignments** on the top bar.
3. Select **Add** and chose **add role assignment**.
   - Select the role created you created in Step 3 and leave "Assign access to" as Active Directory user.
   - Select the active directory created in Step 1 and click **Save**.

CLI:

In the following command, **<role-name>** refers to the name in the sample JSON file, in the earlier example, **Template Deployment**.

```
az ad sp create-for-rbac --name <name-of-service-principal>
--role <role-name>
--output json
```

Find Your Azure Directory Domain Name

For the plugin to provide links to your Azure deployment and Application Insights instance in Azure Portal, you must identify the directory domain for your subscription, as shown below:

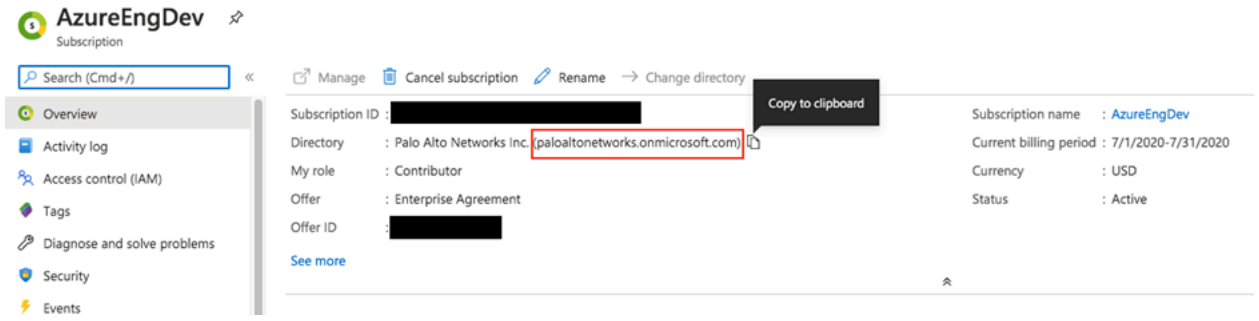

Orchestrate a VM-Series Firewall Deployment in Azure

You can create a maximum of ten orchestrated deployments. Additionally, each orchestrated deployment supports up to 100 applications.

Azure China and Azure Government are not supported.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create a service principal.

Onboard your created service principal credentials to give the Panorama Plugin permissions to make necessary API calls to orchestrate your deployment

1. Select **Setup** > **Service Principal** > **Add**.
2. Enter a **Name** and an optional **Description** to identify the service account.
3. Enter the **Subscription ID** for the Azure subscription you want to monitor.

   You must login to your Azure portal to get this subscription ID.
4. Enter the **Client ID**.The client ID is the Application ID associated with your Azure Active Directory application.
5. Enter the **Client Secret** and re-enter it to confirm.
6. Enter the **Tenant ID**.

   The tenant ID is the Directory ID you saved when you set up the Active Directory application.
7. Click **Validate** to verify that the keys and IDs you entered are valid, and Panorama can communicate with the Azure subscription using the API.

   It can take up to a minute to validate. You can update the page to check your progress.
8. When the service principal is valid, commit your changes.

   The commit ensures the service principal is available when you configure the deployment.

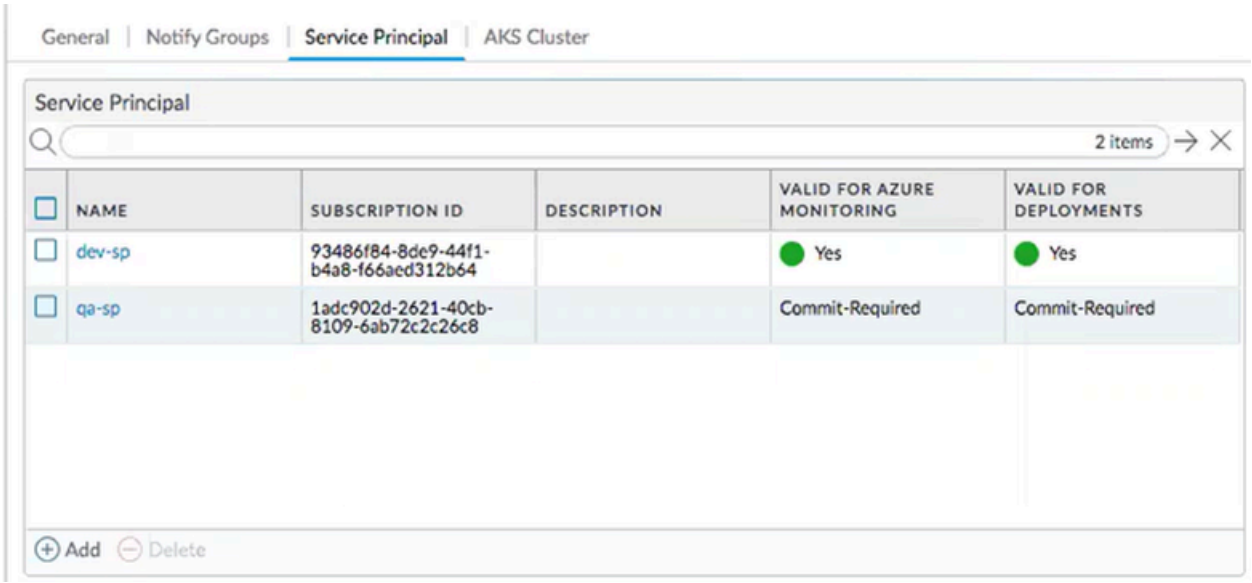

 ©2024 Palo Alto Networks, Inc.

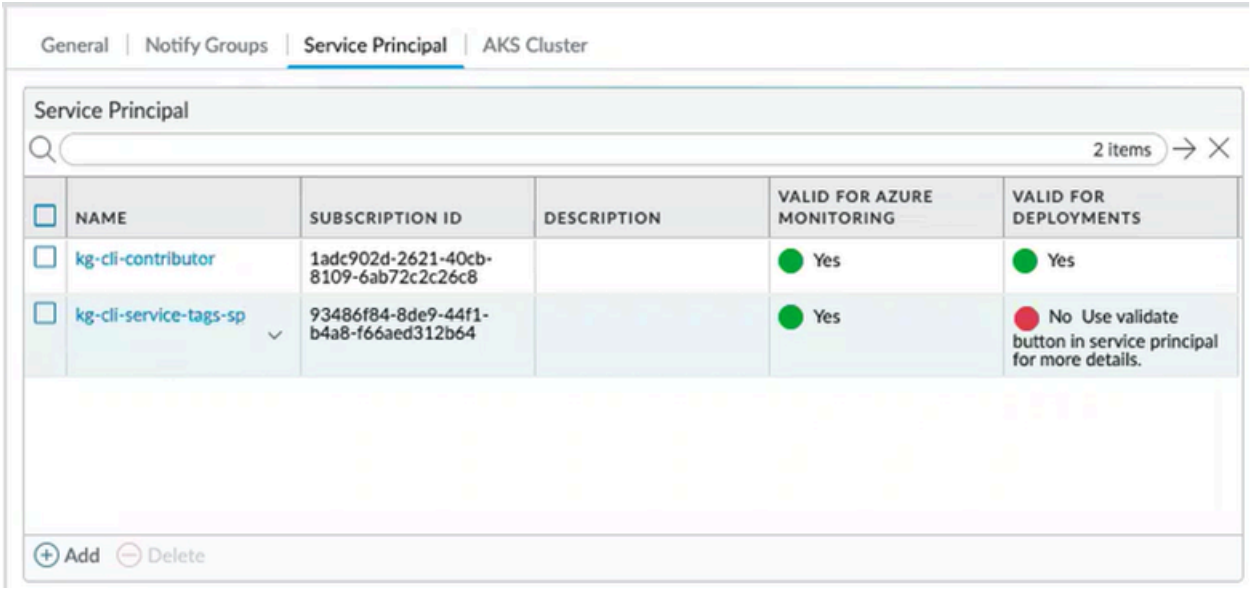

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure your Azure deploymenet.

1. Select **Deployments** and **Add** a configuration.
2. Select **Build** > **General**.
   - Supply a **Name** and an optional **Description**.
   - Choose a service principal from the drop list.

     You must select a valid service principal to enable the **Azure** tab.

     If you don't see your service principal, return to Step 1 and ensure the service principal is valid and committed.

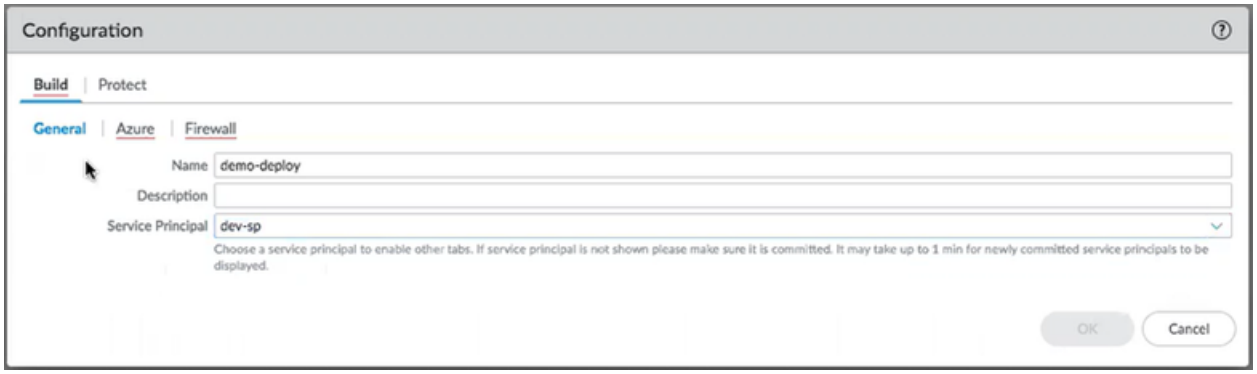

3. On the **Build** > **Azure** tab, select a region.

   The drop list is dynamic—it lists all regions that have a Palo Alto Networks VM-Series Next Generation Firewall image.

   - **Existing VNET**.
     - Select **No** to create a new VNET.

       The plugin uses the VNET CIDR and Directory Domain to create a VNET for you.
     - Select **Yes** to indicate an existing VNET.
   - **VNET CIDR**—Enter your CIDR range. The prefix must be smaller than or equal to /22. For example, 192.168.0.0/22.
   - **Directory Domain**—See Find Your Azure Directory Domain Name. This string is part of the URL for all resources in the subscription, and it helps the plugin link to your deployments.

 ©2024 Palo Alto Networks, Inc.

If you select**Yes** the plugin asks for the VNET Resource Group, the VNET Name, the Security CIDR, and the Directory Domain.

- **VNET Resource Group**—Choose from a list of all resource groups in your selected region.
- **VNET Name**—Choose from a list of VNETS in your chosen resource group.
- **Security CIDR**—Enter your CIDR range. The prefix must be smaller than or equal to /22. For example, 192.168.0.0/22.
- **Directory Domain**—See Find Your Azure Directory Domain Name. This string is part of the URL for all resources in the subscription, and it helps the plugin link to your deployments.

The VNET Resource Group and VNET name help the plugin locate your existing VNET. Anything the plugin deploys goes into a resource group that the plugin manages.

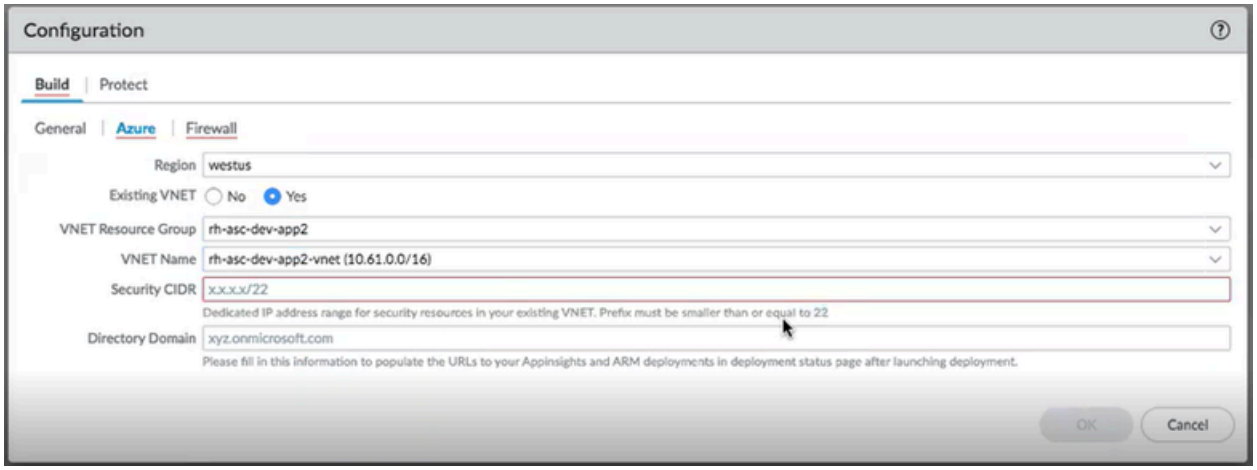

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure the VM-Series firewall stacks for your deployment.

You can deploy the Hub stack to protect Outbound/East-West traffic. You can deploy the Inbound stack to protect inbound traffic. You can also deploy both stacks if all traffic flows need to be protected.

Each inbound stack can secure up to 10 applications.

The configuration parameters are the same for both stacks.

- **License Type**—Select BYOL, Bundle 1, or Bundle 2.
- **License Authcode**—(BYOL only). Enter the authcode sent in your Welcome letter.
- **VM Size**
  - The drop list displays the VM sizes that correlate with the authcode you entered.
  - Bundle1 or Bundle2—Choose any VM size.

Existing Device Group—The device group must be unique across both stacks and deployments. That is, you need a separate dedicated device group for each stack in each deployment.

If you select **No** the plugin creates a device group.

If you select **Yes**, select an existing device group from the dropdown list.

- **Min Firewalls**— A value between 1 and 25 for a VMSS.
- **Max Firewall**— A value between 1 and 25 for a VMSS.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Select **Build** > **Firewall** > **Basic** to configure information common to both Stacks.

For **Image Type**, select **Marketplace Image** or **Custom Image**.

- **Image Resource Group**—(custom image only) Choose the resource group containing your custom image. For a custom image, the list displays all resource groups that contain an image from the region you selected in Step 2.b.
- **Image**—(custom image only) The dropdown list displays all images in your chosen resource group.
- **Software Version**—(Marketplace Image) Only valid software versions are displayed. Consult the Compatibility Matrix for the minimum PAN-OS version.
- **Username**—The administrator user name for the firewall you create. The name must be legal for both VM-Series firewall and Azure. Refer to What are the user name requirments when creating a VM?
- **Password**—The administrator password for the firewall you create. The password must meet the character and length requirements (31 characters) for both VM-Series firewall and Azure. Refer to What are the password requirements when creating a VM?.
- **Confirm Password**—Re-enter your password.
- **Primary Panorama IP**—Specify the Panorama IP address the firewall can use to connect to the Panorama when it boots up. Choose between the public or private IP address displayed in the dropdown list, or type in the Panorama IP address.
- **Secondary Panorama IP**—(Only if Panorama is in HA setup.) Specify the Secondary Panorama IP the firewall can use to connect to the Panorama when it boots up. Choose from dropdown list or type in the correct IP.
- **Configure Device Certificate PIN**. Because these values are encrypted you must enter and confirm each value.
  - **Device Certificate PIN ID**—The device certificate ID.
  - **Confirm Device Certificate PIN ID**
  - **Device Certificate PIN Value**—The certificate PIN value.
  - **Confirm Device Certificate PIN Value**

STEP 5 | Select **Build** > **Firewall** > **Advanced** optional default values.

Check **Advanced** to edit the default values.

- **Autoscaling Metric**—Default is Data Plane CPU Util Percent.
- **Scale In Threshold**—Accept the default or define a scale in thresholdt.
- **Scale Out Threshold**—Accept the default or define a scale in threshold.
- **Jumbo Frame**—Disabled by default.

Click **OK** and commit your changes. Refresh the page until you can see the **Deploy** button, and click **Deploy** to launch the deployment. Once the deployment starts, information is written to the **Deployments** page.

Deployment takes 15-20 minutes to complete.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Select **Azure** > **Deployments** to view deployment status.

- The Resource Group column displays resource groups the plugin has created.
- The firewall's management interface uses the Firewall Access IP to connect to Panorama. You must whitelist this address to ensure that Panorama can connect with Panorama to get the needed configuration.

  If Panorama is deployed in a Public Cloud, make sure to add the Firewall Access IP to the Panorama security group.

  See Ports Used for Panorama to determine which ports you need to open to allow traffic.
- Open the link in the **Deployment Status** column for additional details for each stack.
  - **Hub-Stack**—The Hub stack Public IP matches the Firewall Access IP in the deployment summary because the NAT gateway is the same for egress traffic from the deployment and the management traffic from the firewalls.

    All outbound and East-West traffic should be routed to the **Egress Private IP** for inspection. You can direct traffic to this address if you configured UDRs.
  - **Inbound-Stack**—The Private IP is the address on the Azure internal load balancer that fronts the firewalls. You can direct traffic to this address if you are configuring UDRs.
- Follow the links to view deployment information and Application Insights on Azure.
- The Deployment details can show Success, Warning, and Failure messages

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Configure inbound protection for backend TCP/UDP applications.

The public load balancer that fronts the inbound firewall stack is the entry point for any back-end UDP or TCP applications. Add the following configuration to allow the plugin to manage the necessary load balancer and firewall configuration to route to your backend application.

1. Select **Azure** > **Deployments** and select your deployment.
2. Select the **Protect** tab and click **Add**.
3. Supply the application **Name** and choose a **Protocol**.

   Enter the protection details:

   - **Frontend IP Type**—Select one of New Public IP, Existing Frontend, and Existing Public IP.

     If you select Existing Frontend, the **Frontend Name** lists all known front ends on the load balancer.
   - **Resource Group**—(Existing Public IP only) From the dropdown list, select the resource group where your desired frontend IP address exists.
   - IP Name—(Existing Public IP only) Use to map IP to a frontend on the load balancer, configure the load balancer, and create a NAT rule.
   - **Frontend Port**—Add the frontend port that should be configured to receive traffic on the public load balancer.
   - **Backend IP**—Add the IP address of your backend application.
   - **Backend Port**—Add the port your backend application is expecting to receive traffic on.

   Click **OK**.
4. **Commit** to add the configuration on the load balancer and push to the firewalls.

 ©2024 Palo Alto Networks, Inc.

Create a Custom VM-Series Image for Azure

You can create a custom VM-Series firewall image for later use in your Azure deployment. A custom image gives you the flexibility and consistency to deploy the VM-Series firewall with the PAN-OS version you want to use instead of being restricted to using only an image available though the Azure marketplace. Additionally, your custom image can include the latest content and antivirus updates.

Creating a custom image requires that you remove all private data—user configuration, users, plugin configuration, etc—before creating the VHD. Additionally, Complete the following procedure to prepare and create a custom image.

If the VM-Series firewall used to create your custom image was deployed using a premium disk type, any VM-Series firewall deployed using the custom image must be deployed using the same premium disk type. However, if you create an image using firewall deployed with a standard disk type, you can deploy the firewall using a standard or premium disk type.

STEP 1 | Log in to Azure.

STEP 2 | Deploy the VM-Series firewall from the Azure Marketplace.

STEP 3 | (BYOL license only) Activate your license.

STEP 4 | Upgrade the VM-Series Firewall to PAN-OS 10.0.3. Upgrading to PAN-OS 10.0.3 also upgrades the VM-Series plugin to 2.0.3.

STEP 5 | Access the VM-Series firewall command line interface via SSH using the username and password provided in the Azure Marketplace template.

STEP 6 | Verify that you VM-Series firewall has the correct PAN-OS, VM-Series plugin, content, and antivirus versions.

```
show system info
```

STEP 7 | (BYOL license only) Deactivate your license.

STEP 8 | Perform a private data reset on the VM-Series firewall. This command requires the firewall to reboot. You must wait for the VM-Series firewall to reboot complete before continuing; the reboot can take five to seven minutes.

```
request system private-data-reset
```

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Create a new VHD image from the VM-Series instance.

1. Log in to the Azure CLI.
2. Verify that you are using the correct subscription.

   ```
   az account set --subscription <subscription-id>
   ```
3. Execute the following commands to generalize the VM, allowing it to be imaged for multiple deployments, and create the new VHD.

   ```
   az vm deallocate --resource-group <myResourceGroup> --name
   <myVM>
   ```

   ```
   az vm generalize --resource-group <myResourceGroup> --name
   <myVM>
   ```

STEP 10 | Create a new VM-Series firewall from your custom image.

```
az image create --resource-group <myResourceGroup> --name <myImage>
--source <resource-id-of-VM>
```

STEP 11 | After deploying a VM-Series firewall with your custom image, verify your deployment.

1. You should log in to the firewall using the credentials you used previously.
2. After logging in successfully, verify that your firewall is running the correct PAN-OS version and has the correct content and antivirus versions.

   ```
   show system info
   ```

STEP 12 | (Optional) Copy the custom image to another region.

```
az image copy –source-resource-group <source-rg> –source-object-name
<pa-vm-image-name> –target-location <target-region> –target-resource-
group <destination-rg>
```

 ©2024 Palo Alto Networks, Inc.

Use Azure Security Center Recommendations to Secure Your Workloads

Microsoft has deprecated Azure Security Center support for partner security solutions and replaced it with Azure Sentinel.

When you deploy new workloads within your Azure subscription that is enabled for Azure Security Center, Azure Security Center enables you to secure these workloads in two ways. In one workflow, Azure Security Center recommends you to deploy a new instance of the VM-Series firewall to secure an internet-facing application workload. In the other workflow, Azure Security Center discovers VM-Series firewalls (partner security solutions) that you have deployed within the Azure subscription and you have to then perform additional configuration to connect the VM-Series firewall to Azure Security Center so that you can view alerts on the dashboard. See Azure Security Center Integration for details on the integration and the pros and cons of each workflow:

- Deploy a VM-Series Firewall Based on an Azure Security Center Recommendation
- Connect an Existing VM-Series Firewall From Azure Security Center

Deploy a VM-Series Firewall Based on an Azure Security Center Recommendation

Azure Security Center scans your Azure resources and provides recommendations to secure workloads that need a next-generation firewall. The recommendation displays on the dashboard and you can then either deploy a new instance of the VM-Series firewall from the Azure marketplace or you can use the Azure CLI, Powershell, or an ARM template. The advantage of using a customized deployment using Azure CLI, Powershell, or an ARM template is that you can deploy the VM-Series firewall within the same resource group as the workload that you need to secure. When you deploy the VM-Series firewall using the Azure marketplace, Azure requires that you deploy the firewall into a new resource group or an empty resource group only. Therefore, the marketplace deployment requires you to then ensure that the traffic from the workload you want to secure is steered to the firewall that is in a different resoource group.

STEP 1 | Log in to your Azure portal and access the Security Center dashboard.

STEP 2 | Select **Recommendations**.

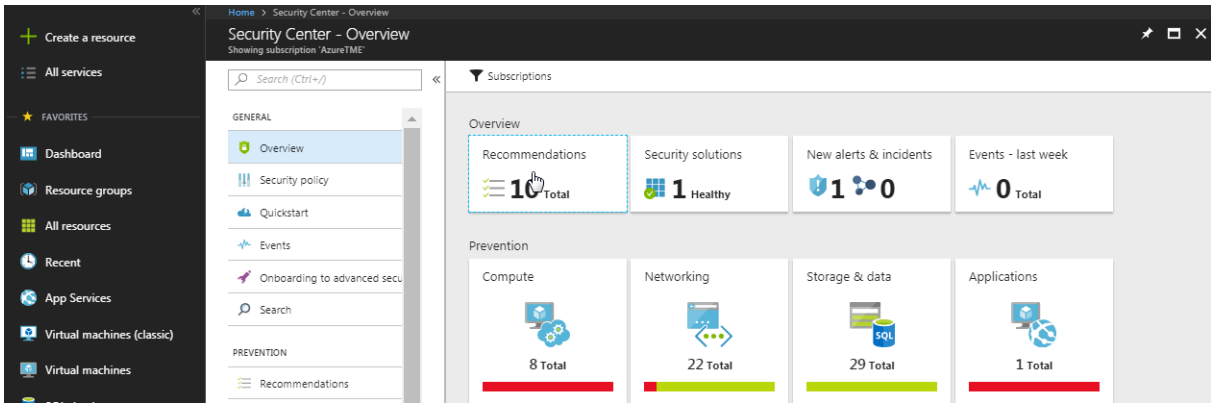

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Select **Add a Next Generation Firewall**, select the workload you want to secure.

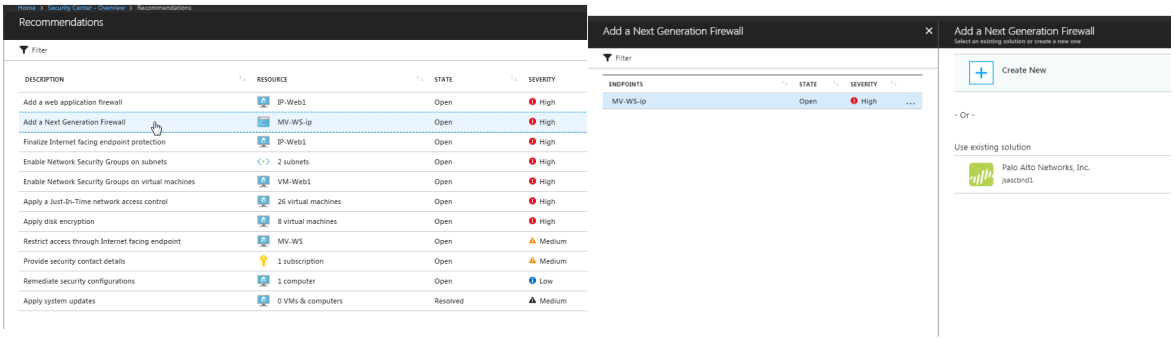

STEP 4 | Choose whether you want to deploy a new instance of the VM-Series firewall or use an existing instance of the VM-Series firewall.

To use this workflow, stage a workload with a public IP address that is exposed to the internet and deploy an instance of the VM-Series firewall in a new resource group. Then, delete the workload you staged, and deploy your production workloads within the resource group in which you deployed the VM-Series firewall.

- To **Create New**, see Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template).
- To **Use existing solution**, select the VM-Series firewall that you have previously deployed.

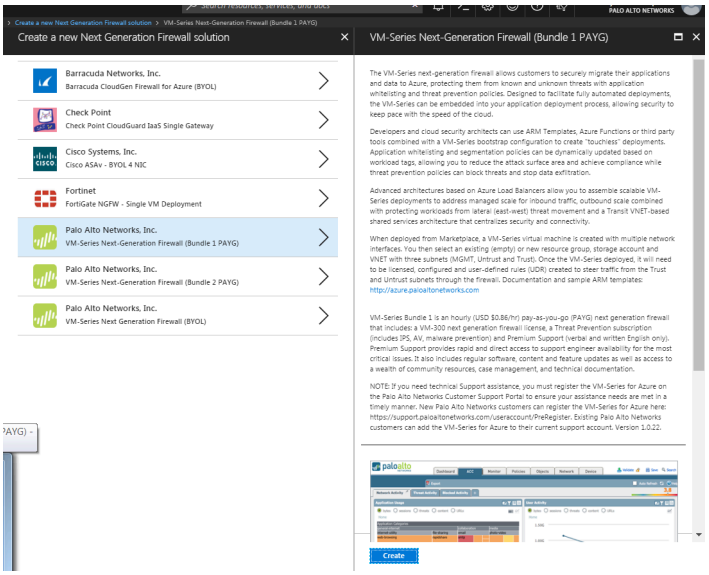

Connect an Existing VM-Series Firewall From Azure Security Center

When Azure Security Center detects that you have deployed the VM-Series firewall within the Azure subscription, it displays the firewall as a security solution. You can then connect the VM-Series firewall to Security Center using the Common Event Format (CEF) over Syslog, and view firewall logs as alerts on the Security Center dashboard.

STEP 1 | Log in to your Azure portal and access the Security Center dashboard.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Select Security Solutions to view all available VM-Series firewalls within this Azure subscription.

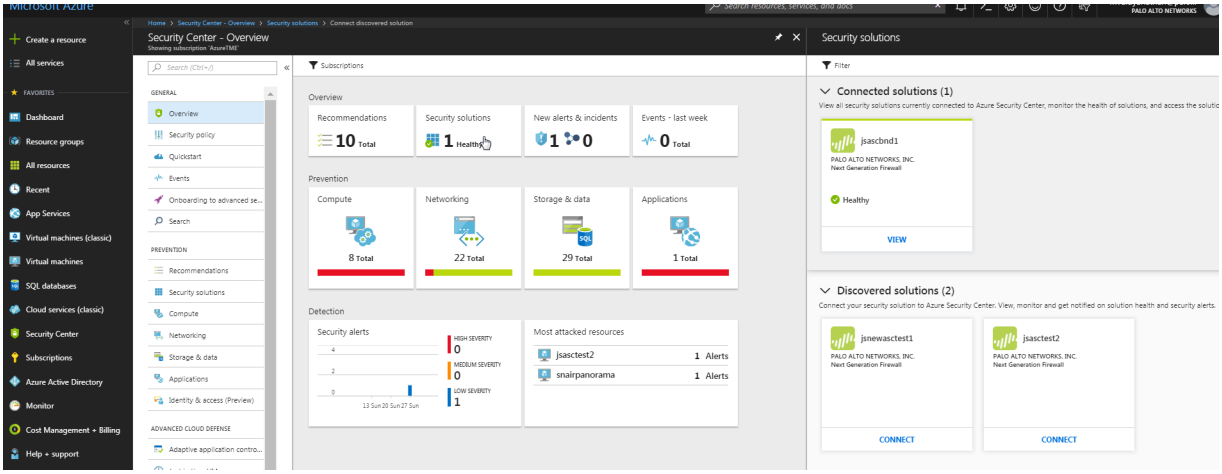

STEP 3 | Expand Discovered solutions, and select the VM-Series firewall instance that is in the same resource group as the workload you want to secure and click **Connect**.

To view firewall logs as alerts on the Security Center dashboard, you need to follow the four-step process that displays on screen.

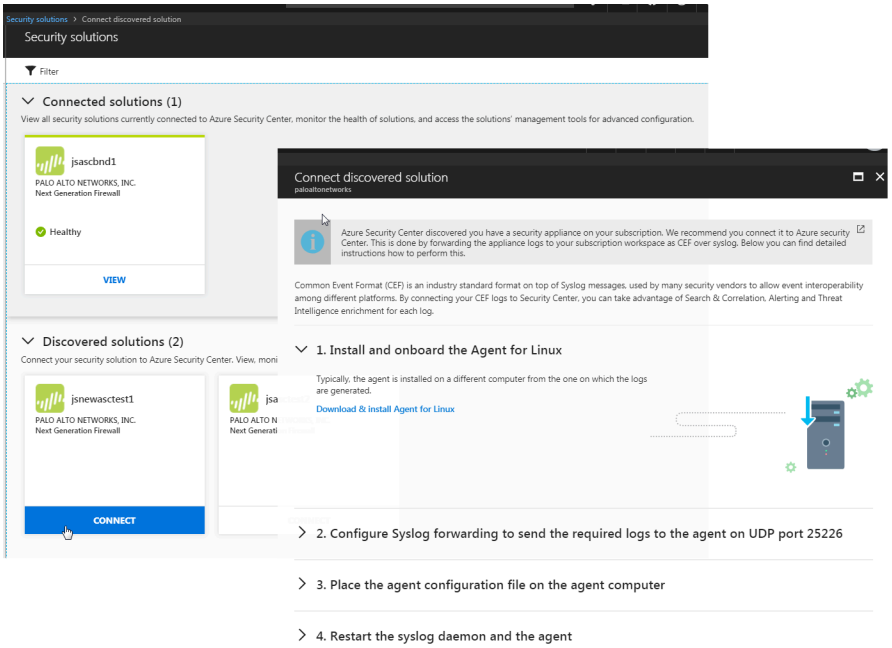

 ©2024 Palo Alto Networks, Inc.

STEP 4 | On successfully connecting the VM-Series firewall to Security Center, the VM-Series firewall displays in the Connected solutions list.

Click View to verify that the firewall is protecting the workload that you need to secure.

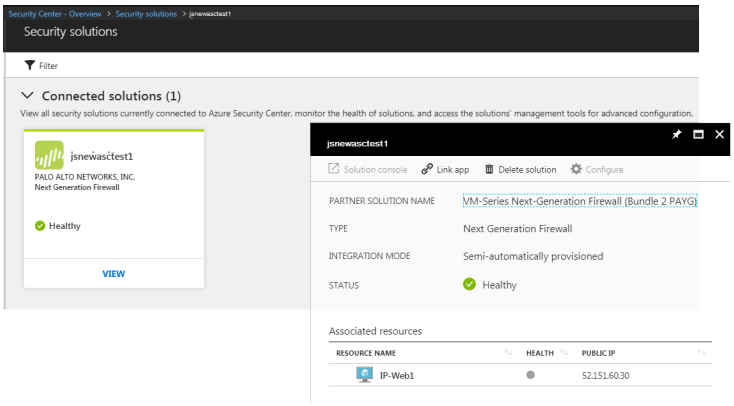

Use Panorama to Forward Logs to Azure Security Center

If you are using Panorama to manage your firewalls, you can use templates and device groups to forward firewall logs to Azure Security Center. With the default Azure Security Center Log Forwarding profile, Threat and WildFire Submissions logs of low, medium, high, or critical severity generated on the firewall are displayed as security alerts on the Azure Security Center dashboard. So that you can focus and triage alerts more efficiently, you can set up granular log filters to only forward logs of interest to you, or forward high and critical severity logs only. You can also selectively attach the log forwarding profile to a few Security policy rules based on your applications and security needs.

To enable the Azure Security Center integration from Panorama, use the following workflow.

STEP 1 | Add the firewall as a managed device on Panorama.

STEP 2 | From Panorama, create a template and a device group to push log forwarding settings to the firewalls that will be forwarding logs to Azure Security Center.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Specify the log types to forward to the Logging Service.

The way you enable forwarding depends on the log type. For logs that are generated based on a policy match, you use a log forwarding profile within a device group, and for other logs types you use the Log Settings configuration within a template.

1. Configure forwarding of System, Configuration, User-ID, and HIP Match logs.
   1. Select **Device** > **Log Settings**.
   2. Select the **Template** that contains the firewalls you want to forward logs to the Logging Service.
   3. For each log type that you to forward to the Logging Service, **Add** a match list filter. Give it a **Name**, optionally define a **Filter**.
   4. **Add** Built-in Actions and enter a **Name**. The Azure-Security-Center-Integration action will be auto selected. Click **OK**.

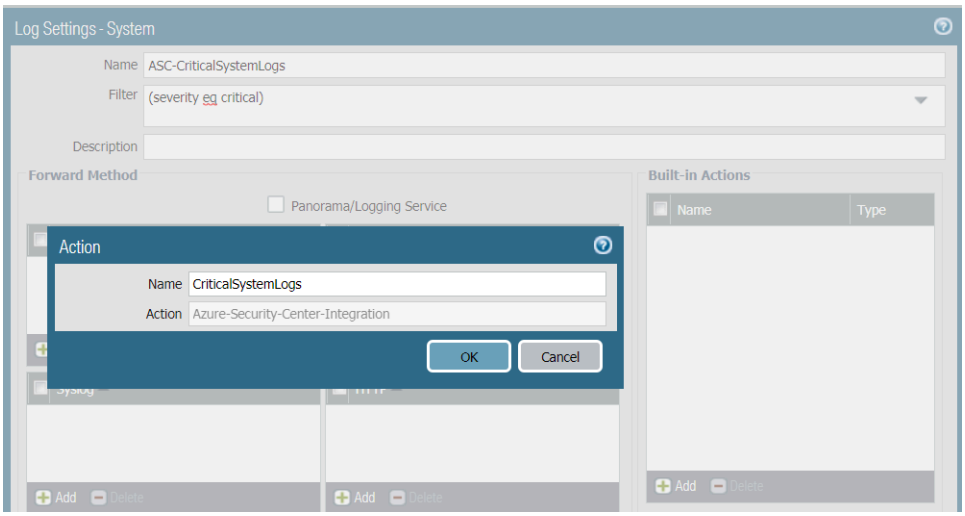

   5. Click **OK**.
2. Configure forwarding of all other log types that are generated when a policy match occurs such as Traffic, Threat, WildFire Submission, URL Filtering, Data Filtering, and

 ©2024 Palo Alto Networks, Inc.

Authentication logs. To forward these logs, you must create and attach a log forwarding profile to each policy rule for which you want to forward logs.

1. Select the **Device Group**, and then select **Objects** > **Log Forwarding** to **Add** a profile. In the log forwarding profile match list, add each log type that you want to forward.
2. Select **Add** in Built-in Actions to enable the firewalls in the device group to forward the logs to Azure Security Center.

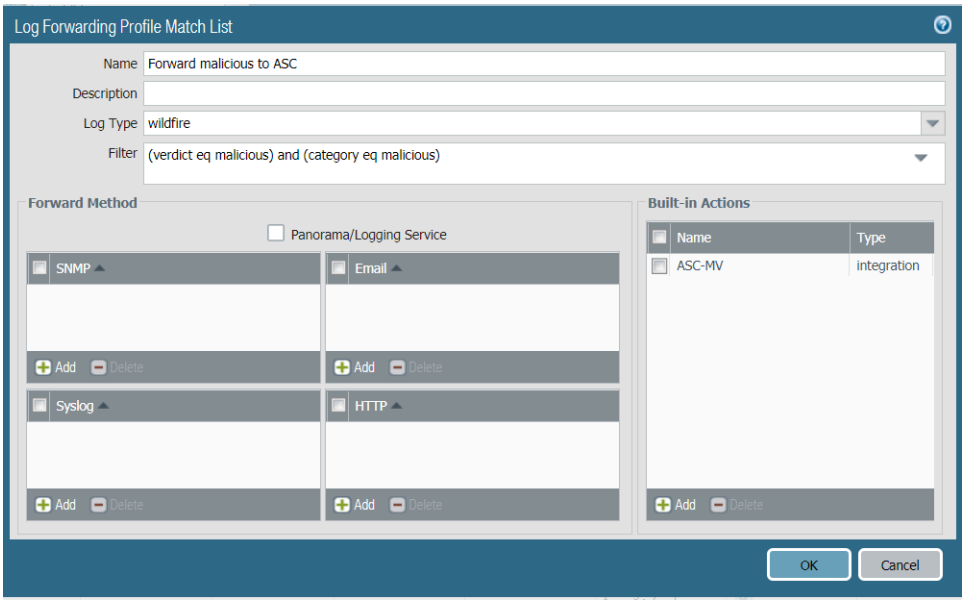

3. Create basic security policy rules in the device group you just created and select **Actions** to attach the Log Forwarding profile you created for forwarding logs to Azure Security Center. Until the firewall has interfaces and zones and a basic security policy, it will not let any traffic through, and only traffic that matches a security policy rule will be logged (by default).
4. For each rule you create, select **Actions** and select the Log Forwarding profile that allows the firewall to forward logs to Azure Security Center.

STEP 4 | Commit your changes to Panorama and push them to the template and device group you created.

STEP 5 | Verify that the firewall logs are being forwarded to Azure Security Center.

1. Log in the Azure portal, select **Azure Security Center**.
2. Verify that you can see firewall logs as Security alerts on the Azure Security Center dashboard.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on Azure Stack

You can deploy the VM-Series firewall on Azure Stack to secure inter-subnet traffic between applications in a multi-tier architecture and outbound traffic from servers within your Azure Stack deployment. If you want to use the VM-Series firewall as a gateway that secures inbound traffic destined to the servers within your Azure Stack deployment, you must deploy a NAT appliance in front of the firewall that receives inbound traffic and forwards it to the firewall. The NAT appliance is required because on Azure Stack you cannot assign a public IP address to a non-primary interface of a virtual machine, such as the VM-Series firewall.

The VM-Series firewall on Azure stack does not have support for bootstrapping, Azure Application Insights, or the Azure Security Center integration.

Unlike on public Azure, you do not have a solution template to deploy the VM-Series firewall on Azure Stack. Therefore, you must use an ARM template to deploy the VM-Series firewall. To get started, you can use the community supported sample ARM template on GitHub, and then develop your own ARM template for production deployments.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Download marketplace items from Azure to AzureStack.

To deploy the VM-Series firewall on Azure Stack, you need access to the BYOL offer of the VM-Series firewall PAN-OS image (8.1 or later). You can download the image directly from the Azure Marketplace to Azure Stack in a connected deployment.

STEP 2 | Access the Azure Stack portal.

Your Azure Stack operator (either a service provider or an administrator in your organization), should provide the correct URL to access the portal.

STEP 3 | Deploy the VM-Series firewall.

A solution template for the VM-Series firewall is not available on Azure Stack. Therefore, you must reference the image that you downloaded in the previous step, in an ARM template to deploy the VM-Series firewall. To get started, you can deploy the sample ARM template that is available on GitHub under the community supported policy:

1. Get the sample Azure Stack GitHub template.
   - Select azurestackdeploy.json to view the contents.
   - Click Raw and copy the contents of the JSON file.
2. Deploy the sample GitHub template.

   You can deploy the firewall in a existing resource group that is empty or into a new resource group. The default VNet in the template is 192.168.0.0/16, and it deploys a VM-Series firewall has three network interfaces, one management interface

 ©2024 Palo Alto Networks, Inc.

on 192.168.0.0/24 subnet and two dataplane interfaces on 192.168.1.0/24 and 192.168.2.0/24 subnets. You can customize these subnets to match your needs.

- Log in to the Azure Stack portal.
- Select **New** > **Custom** > **Template deployment**.

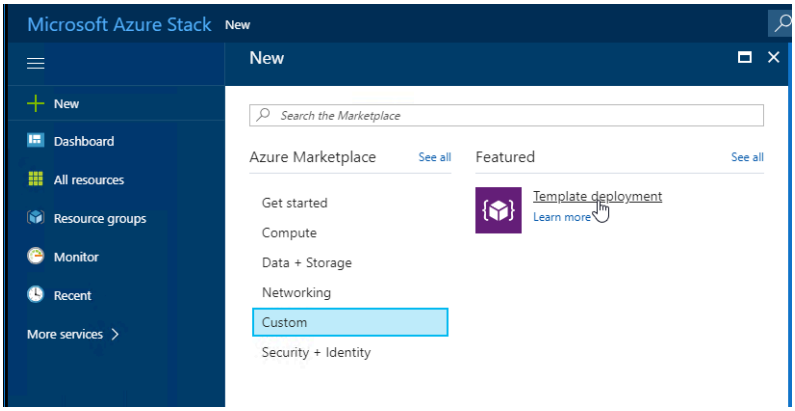

- **Edit template**, delete all existing content in the template, and paste the JSON template contents you copied earlier and**Save**.

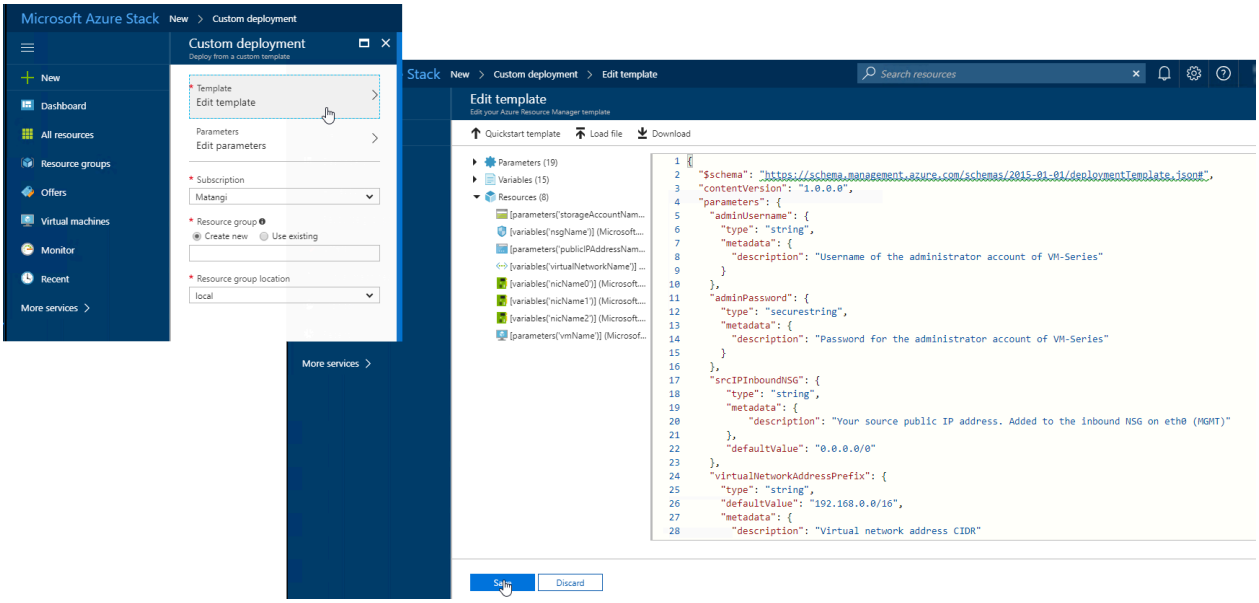

- **Edit parameters**, enter the values for the required parameters and modify the defaults if you need to, then click **OK**.

 ©2024 Palo Alto Networks, Inc.

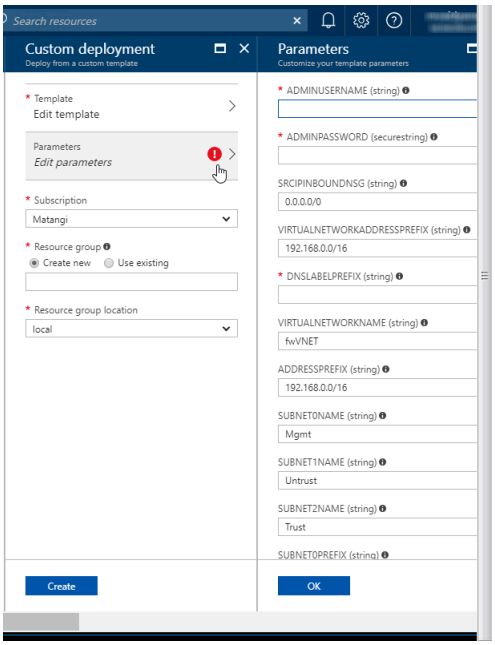

- Choose the **Subscription** you want to use, and then click **OK**.
- Choose an existing **Resource Group** that is empty or create a new one, and click **OK**.
- Click **Create**. A new tile on the dashboard displays the progress of the template deployment.

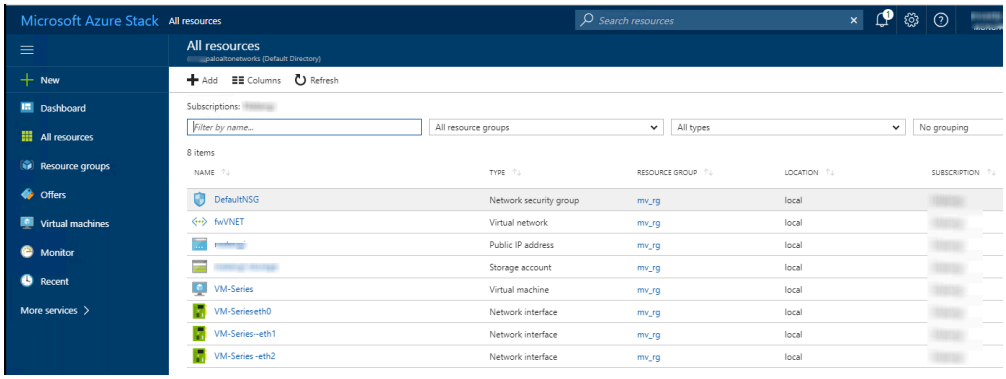

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Next Steps:

1. Log in to the web interface of the firewall.

   Using a secure connection (https) from your web browser, log in to the DNS name for the firewall. Enter the username/password you defined earlier. You will see a certificate warning; that is okay. Continue to the web page.

2. Activate the licenses on the VM-Series firewall.

   1. Create a Support Accountand Register the VM-Series Firewall (with auth code)
   2. On the firewall web interface, select **Device** > **Licenses** and select **Activate feature using authentication code**.
   3. Enter the capacity auth-code that you registered on the support portal. The firewall will connect to the update server (updates.paloaltonetworks.com), and download the license and reboot automatically.
   4. Log back in to the web interface on the **Dashboard**, confirm that a valid **Serial#** displays.

      The **VM Mode** displays as Microsoft Azure.

      If the term Unknown displays, it means the device is not licensed. To view traffic logs on the firewall, you must install a valid capacity license.

STEP 5 | Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

 ©2024 Palo Alto Networks, Inc.

Enable Azure Application Insights on the VM-Series Firewall

The VM-Series firewall on Azure can publish custom PAN-OS metrics natively to Azure Application Insights that you can use to monitor the firewalls directly from the Azure portal. These metrics allow you to assess performance and usage patterns that you can use to set alarms and take action to automate events such as launching or terminating instances of the VM-Series firewalls. See Custom PAN-OS Metrics Published for Monitoring for a description on the metrics that are available.

STEP 1 | On the Azure portal, create your Application Insights instance to monitor the firewall and copy the **Instrumentation Key** from **Configure** > **Properties**.

The firewall needs this key to authenticate to the Application Insights instance and publish metrics to it. See VM-Series on Azure Service Principal Permissions for the permissions required.

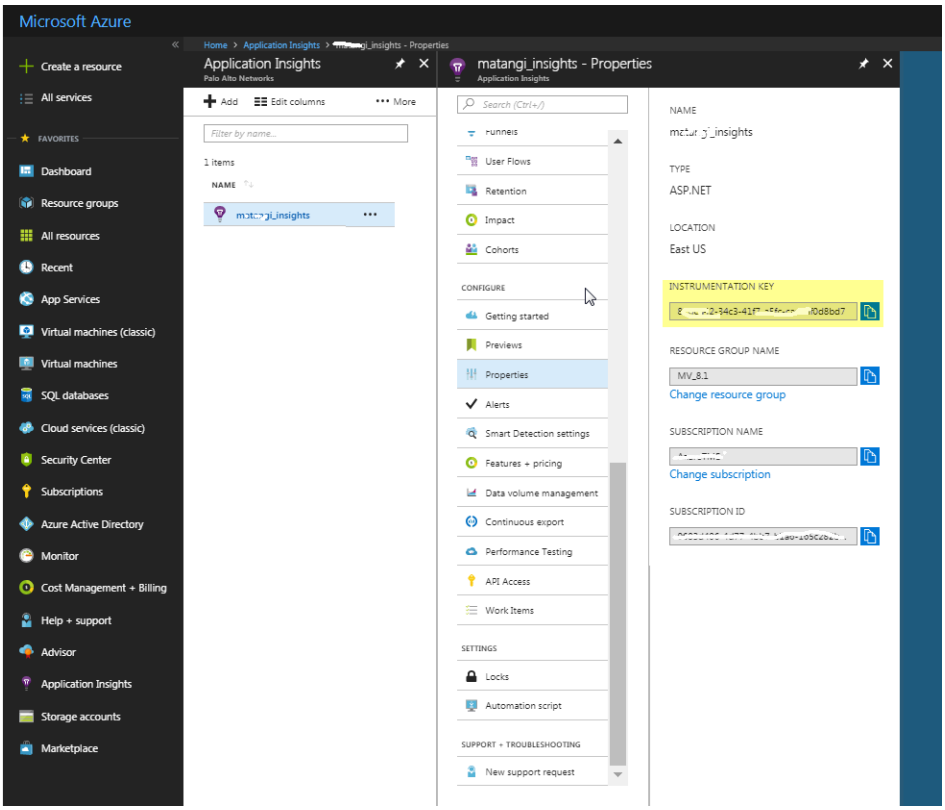

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enable the firewall to publish metrics to your Application Insights instance.

1. Log in to the VM-Series firewall on Azure.
2. Select **Device** > **VM-Series** > **Azure**.
3. Edit **Azure Application Insights** and enter the Instrumentation Key you copied earlier.

   The default interval for publishing metrics is five minutes. You can change this to vary from 1-60 minutes.

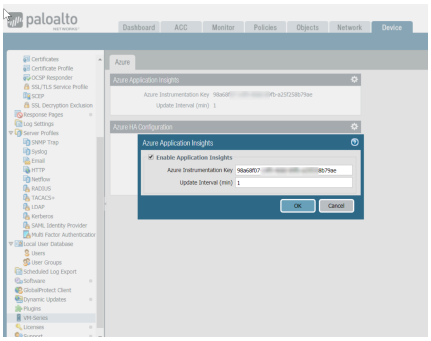

4. **Commit** your changes.

   The firewall generates a system log to it record the success or failure to authenticate to Azure Application Insights.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Verify that you can view the metrics on the Azure Application Insights dashboard.

1. On the Azure portal, select the Application Insights instance, and select **Monitoring** > **Metrics** to view the PAN-OS custom metrics.

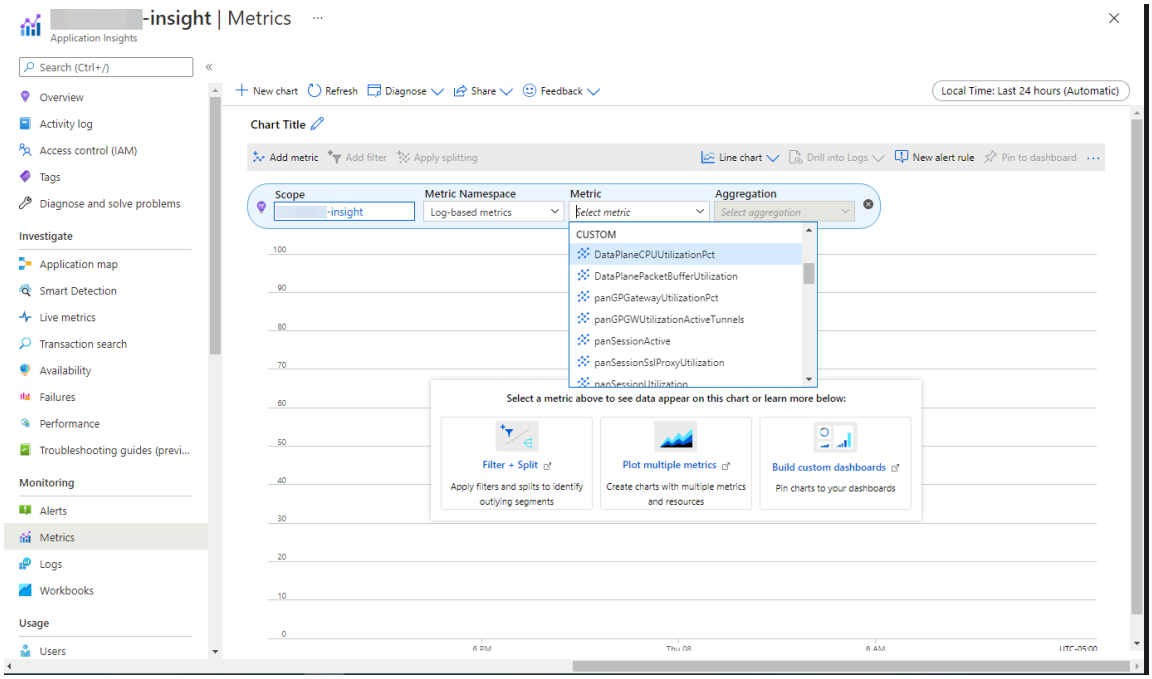

2. Select the metric(s) that you want to monitor for trends and trigger alerts. Refer to the Microsoft Azure documentation for details on exploring metrics on Application Insights.

 ©2024 Palo Alto Networks, Inc.

Monitoring on Azure

Monitoring on Microsoft® Azure® enables you to dynamically update security policy rules to consistently enforce Security policy across all assets deployed within your Azure subscription. To enable this capability, you need to install the Panorama plugin for Azure and enable API communication between Panorama and your Azure subscriptions. Panorama can then collect the IP address-to-tag mapping for all your Azure assets and push or distribute Azure resources information to your Palo Alto Networks® firewall(s).

- About Monitoring on Azure
- Set Up the Azure Plugin for Monitoring on Panorama
- Attributes Monitored Using the Panorama Plugin on Azure

About Monitoring on Azure

As you deploy or terminate virtual machines in the Azure public cloud, you can use the Panorama plugin for Azure to consistently enforce security policy rules on these workloads.

The Panorama plugin for Azure is built for scale and allows you to monitor up to 100 Azure subscriptions on the Azure public cloud. With this plugin, you use Panorama as an anchor to poll your subscriptions for tags, and then distribute the metadata (IP address-to-tag mapping) to many firewalls in a device group. Because Panorama communicates with your Azure subscriptions to retrieve Azure resource information, you're able to streamline the number of API calls made to the cloud environment. Although you can define Security policy locally on the firewall, using Panorama and the plugin centralizes Security policy management, ensuring consistent policies for hybrid and cloud-native architectures.

See the Panorama plugin version information in the Compatibility Matrix.

Set Up the Azure Plugin for Monitoring on Panorama

To find all the workloads that your organization has deployed in the Azure cloud, you need to install the Azure plugin on Panorama and configure *Monitoring Definitions* that enable Panorama to authenticate to your Azure subscription(s) and retrieve information on the Azure workloads. Panorama retrieves the primary private IP address of Azure resources and the associated tags. For a list of the metadata elements that Panorama supports, see Attributes Monitored Using the Panorama Plugin on Azure.

After Panorama fetches the attributes, to push the resource information from Panorama to the firewalls, you must add the firewalls (hardware or VM-Series) as managed devices on Panorama, and group the firewalls into one or more Device Groups. You can then specify which device groups are part of the *Notify Group*, which is a configuration element in a Monitoring Definition, that Panorama uses to register the IP address-to-tag mapping it retrieves from Azure.

Finally, to consistently enforce Security policies across your Azure workloads, you must set up Dynamic Address Groups and reference them in policy rules that allow or deny traffic to the IP addresses of the Azure resources. For streamlining your configuration and managing policies and objects centrally from Panorama, you can define the Dynamic Address Groups and Security policy rules on Panorama and push them to the firewalls instead of managing the Dynamic Address Groups and Security policy rules locally on each firewall.

 ©2024 Palo Alto Networks, Inc.

The Azure plugin is for monitoring Azure resources on the Azure public cloud. Azure Government or Azure China are not supported.

- Planning Checklist for VM Monitoring with the Azure Plugin
- Install the Azure Plugin
- Configure the Azure Plugin for Monitoring

Planning Checklist for Monitoring with the Azure Plugin

- Set up the Active Directory application and a Service Principal to enable API access—For Panorama to interact with the Azure APIs and collect information on your workloads, you need to create an Azure Active Directory Service Principal. This Service Principal has the permissions required to authenticate to the Azure AD and access the resources within your subscription.

  To complete this set up, you must have permissions to register an application with your Azure AD tenant, and assign the application to a role in your subscription. If you don't have the necessary permissions, ask your Azure AD or subscription administrator to create a Service Principal with an IAM role of reader or the custom permissions specified in VM-Series on Azure Service Principal Permissions.
- Make sure that the subscription ID is unique across Service Principals. Panorama allows you to use only one service principal to monitor an Azure subscription. You can monitor up to 100 Azure subscriptions, with 100 Service principal resources. Starting with Panorama plugin for Azure version 3.2.0, you can monitor up to 500 Azure subscriptions. Take note of the processing times described in the table below.

	Number of Subscriptions				
	100	200	300	400	500
System Resources Utilization	• CPU: 22% • Memory: 0.025 MB	• CPU: 23.8% • Memory: 0.025 MB	• CPU: 31.8% • Memory: 0.025 MB	• CPU: 28.6% • Memory: 0.025 MB	• CPU: 39.2% • Memory: 0.025 MB
Average Time to Process All Monitoring Definitions	1 Hour, 15 Minutes	2 Hours, 30 Minutes	3 Hours, 30 Minutes	4 Hours	5 Hours
Average Tag Update Processing Time	5 minutes	10 minutes	15 minutes	20 minutes	25 minutes

The information in the table above was captured on an instance with 8 vCPUs and 32GB of memory.

 ©2024 Palo Alto Networks, Inc.

- Panorama can push up to 8000 IP address-to tag mappings to the firewalls or virtual system assigned to a device group. Review the requirements for Panorama and the managed firewalls:
  - Minimum system requirements (see the Panorama Plugin information in the Compatibility Matrix):
  - You must add the firewalls as managed devices on Panorama and create Device Groups so that you can configure Panorama to notify these groups with the information it retrieves. Device groups can include VM-Series firewalls or virtual systems on the hardware firewalls.
  - The number of tags that the Panorama plugin can retrieve and register is as follows:

    On Panorama running 8.1.3 or later managing firewalls running PAN-OS 8.1.3 or lower, the firewalls or virtual systems included within a device group can have 7000 IP addresses with 10 tags each, or 6500 IP addresses with 15 tags each.

    On Panorama 8.1.3 or later managing firewalls running PAN-OS 8.0.x, 2500 IP addresses with 10 tags each.

    - If your Panorama appliances are in a high availability configuration, you must manually install the same version of the Azure plugin on both Panorama peers.

      You configure the Azure plugin on the active Panorama peer only. On commit, the configuration is synced to the passive Panorama peer. Only the active Panorama peer polls the Azure subscriptions you have configured for Monitoring.

Install the Azure Plugin

To get started with Monitoring on Azure, you need to download and install the Azure plugin on Panorama. If you have a Panorama HA configuration, repeat this installation process on each Panorama peer.

If you currently have installed a Panorama plugin, the process of installing (or uninstalling) another plugin requires a Panorama reboot to enable you to commit changes. So, install additional plugins during a planned maintenance window to allow for a reboot.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

 ©2024 Palo Alto Networks, Inc.

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Log in to the Panorama Web Interface, select **Panorama** > **Plugins** and click **Check Now** to get the list of available plugins.

STEP 2 | Select **Download** and **Install** the plugin.

After you successfully install, Panorama refreshes and the Azure plugin displays on the **Panorama** tab.

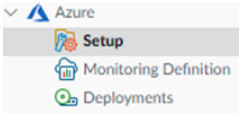

STEP 3 | Restart Panorama.

Select **Panorama** > **Setup** > **Operations** > **Reboot Panorama**

Configure the Azure Plugin for Monitoring

To begin monitoring the resources in your Azure public cloud deployment, after you Install the Azure Plugin you must create a Monitoring Definition. This definition specifies the Service Principal that is authorized to access the resources within the Azure subscription you want to monitor and the Notify Group that includes the firewalls to which Panorama should push all the IP-address-to-tag mappings it retrieves. In order to enforce policy, you must then create Dynamic Address Groups and reference them in Security policy. The Dynamic Address Groups enable you to filter the tags you want to match on, so that the firewall can get the primary private IP address registered for the tags, and then allow or deny access to traffic to and from the workloads based on the policy rules you define.

STEP 1 | Log in to the Panorama web interface.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Set up the following objects for enabling Monitoring on Azure.

- Add a Service Principal.

  The Service Principal is the service account that you created on the Azure portal. This account is attached to the Azure AD and has limited permissions to access and monitor the resources in your Azure subscription.

  1. Select **Panorama** > **Plugins** > **Azure** > **Setup** > **Service Principal** > **Add**.

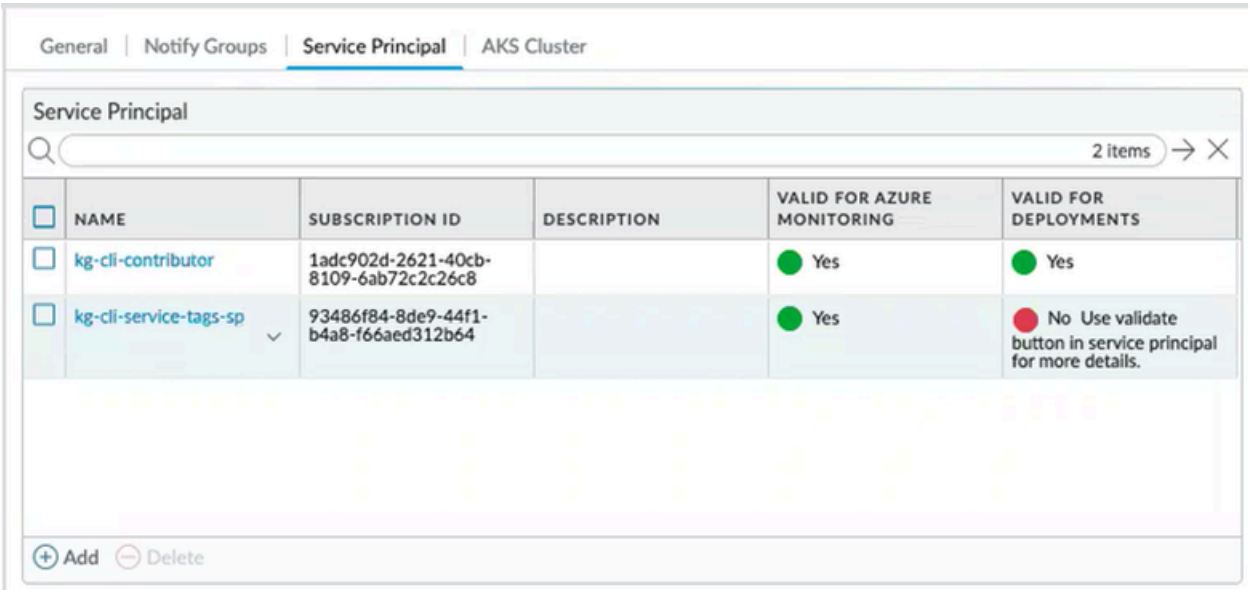

  2. Enter a **Name** and optionally a **Description** to identify the service account.
  3. Enter the **Subscription ID** for the Azure subscription you want to monitor. You must login to your Azure portal to get this subscription ID.
  4. Enter the **Client Secret** and re-enter it to confirm.
  5. Enter the **Tenant ID**. The tenant ID is the Directory ID you saved when you set up the Active Directory application.
  6. Click **Validate** to verify that the keys and IDs you entered are valid ,and Panorama can communicate with the Azure subscription using the API.

- Add a notify group.

  1. Select **Panorama** > **Plugins** > **Azure** > **Setup** > **Notify Groups** > **Add**.

 ©2024 Palo Alto Networks, Inc.

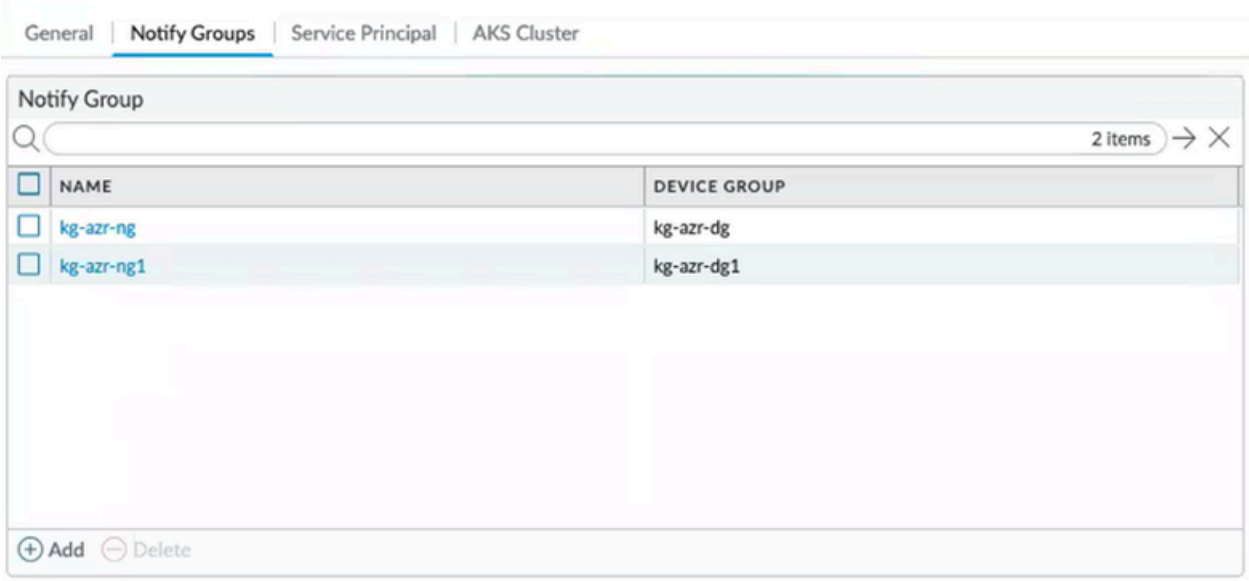

2. Enter a **Name** and optionally a **Description** to identify the group of firewalls to which Panorama pushes the information it retrieves.
3. Select the **Device Groups**, which are a group of firewalls or virtual systems, to which Panorama will push the information (IP address-to-tag mapping) it retrieves from your Azure subscriptions. The firewalls use the update to determine the most current list of members that constitute dynamic address groups referenced in policy.

   Think through your Device Groups carefully.

   - Because a Monitoring Definition can include only one notify group, make sure to select all the relevant Device Groups within your notify group. If you want to deregister the tags that Panorama has pushed to a firewall included in a notify group, you must delete the Monitoring Definition.
   - To register tags to all virtual systems on a firewall enabled for multiple virtual systems, you must add each virtual system to a separate device group on Panorama and assign the device groups to the notify group. Panorama will register tags to only one virtual system, if you assign all the virtual systems to one device group.
4. Verify that monitoring is enabled on the plugin. This setting must be enabled for Panorama to communicate with the Azure public cloud for Monitoring.

   The checkbox for **Enable Monitoring** is on **Panorama** > **Plugins** > **Azure** > **Setup** > **General**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create a **Monitoring Definition**.

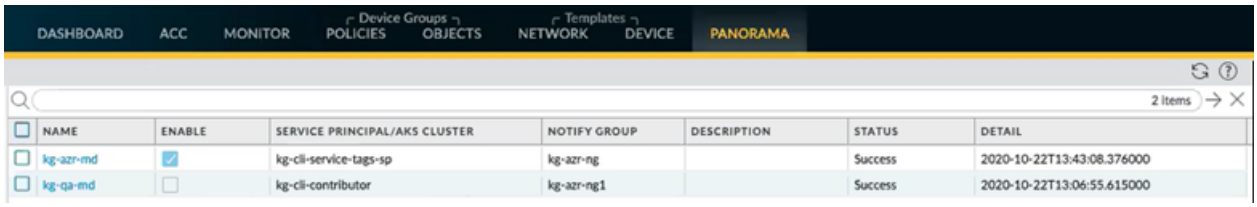

When you add a new Monitoring definition, it is enabled by default.

- Select **Panorama** > **Plugins** > **Azure** > **Monitoring Definition**, to **Add** a new definition.
- Enter a **Name** and optionally a **Description** to identify the Azure subscription for which you use this definition.
- Select the **Service Principal** and **Notify Group**.

  Panorama requires the keys and IDs that you specify in the Service Principal configuration to generate an Azure Bearer Token which is used in the header of the API call to collect information on your workloads.

STEP 4 | **Commit** the changes on Panorama.

Verify that the status for the Monitoring Definition displays as Success. If it fails, verify that you entered the Azure Subscription ID accurately and provided the correct keys and IDs for the Service Principal.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Verify that you can view the information on Panorama, and define the match criteria for Dynamic Address Groups.

Some browser extensions may block API calls between Panorama and Azure which prevents Panorama from receiving match criteria. If Panorama displays no match criteria and you are using browser extensions, disable the extensions and Synchronize Dynamic Objects to populate the tags available to Panorama.

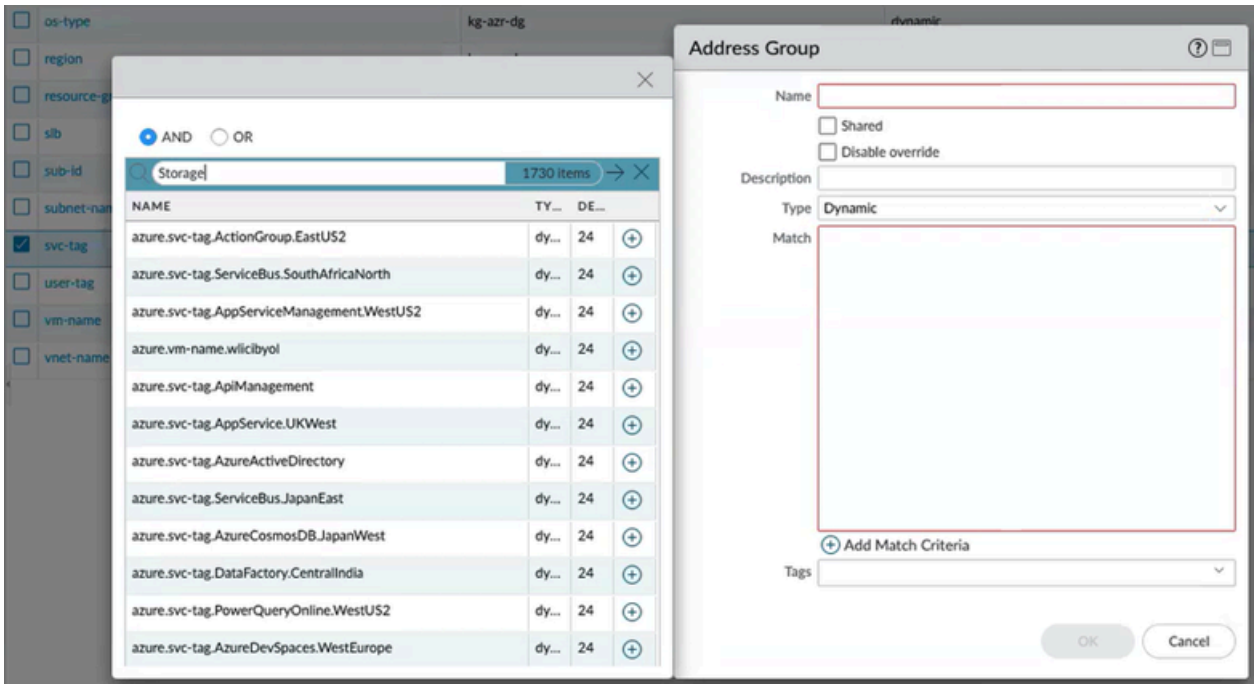

On HA failover, the newly active Panorama attempts to reconnect to the Azure cloud and retrieve tags for all monitoring definitions. If there is an error with reconnecting even one monitoring definition, Panorama generates a system log message

```
Unable to process subscriptions after HA switch-over;
 user-intervention required.
```

When you see this error, you must log in to Panorama and fix the issue, for example remove an invalid subscription or provide valid credentials, and commit your changes to enable Panorama to reconnect and retrieve the tags for all monitoring definitions. Even when Panorama is disconnected from the Azure cloud, the firewalls have the list of all tags that had been retrieved before failover, and can continue to enforce policy on that list of IP addresses. Panorama removes all tags associated with the subscription only when you delete a monitoring definition. As a best practice, to monitor this issue, configure action-oriented log forwarding to an HTTPS destination *from Panorama so that you can take immediate action.*

 ©2024 Palo Alto Networks, Inc.

Attributes Monitored Using the Panorama Plugin on Azure

When using the Panorama plugin for Azure, Panorama gathers the following set of metadata elements or attributes on the virtual machines in your Microsoft® Azure® deployment. Panorama can retrieve a total of 32 tags for each VM, 11 predefined tags and up to 21 user-defined tags.

The maximum length of a tag can be 127 characters. If a tag is longer than 127 characters, Panorama does not retrieve the tag and register it on the firewalls. Also the tags should not include non-ASCII special characters such as { or ".

The following attributes are monitored in all Panorama plugin for Azure versions.

Virtual Machine

VM Monitoring	Example
VM Name	azure-tag.vm-name.web_server1
Network Security Group Name	azure-tag.nsg-name.myNSG
OS Type	azure-tag.os-type.Linux
OS Publisher	azure-tag.os-publisher.Canonical
OS Offer	azure-tag.os-offer.UbuntuServer
OS SKU	azure-tag.os-sku.14.04.5-LTS
Subnet	azure-tag.subnet.webtier
VNet	azure-tag.vnet.untrustnet
Azure Region	azure-tag.region.east-us
Resource Group Name	azure-tag.resource-group.myResourceGroup
Subscription ID	azure.sub-id.93486f84-8de9-44f1-b4a8-f66aed312b64
User Defined Tags Up to a maximum of 21 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 21 tags are available for use on Panorama and the firewalls.	azure-tag.mytag.value

Load Balancer

 ©2024 Palo Alto Networks, Inc.

Panorama plugin on Azure version 3.0 or later supports tags for each application gateway and standard load balancer (both public and private IP addresses). Each load balancer has predefined tags for resource group, load balancer name and region, and supports up to 21 user-defined tags specific to load balancing.

Load Balancer Tags	Example
Load Balancer	azure.<type>.myLoadBalancer
Azure Region	azure-tag.region.east-us
Resource Group Name	azure-tag.resource-group.myResourceGroup
User Defined Tags Up to a maximum of 21 user defined tags are supported. The user-defined tags are sorted alphabetically, and the first 21 tags are available for use on Panorama and the firewalls.	azure-tag.mytag.value

Subnet/VNET

Panorama plugin on Azure version 3.0 or later supports tags for each Subnet and VNET in your subscription. Each subnet and VNET tag is associated with the full IP CIDR range so you can create policies based on a CIDR range rather than individual IP addresses. The plugin queries every subnet and VNET in your subscription and creates tags for them.

Subnet and VNET Tags	Example
Subnet Name	azure.subnet-name.web
VNET Name	azure.vnet-name.myvnet

Service Tag Monitoring

Panorama plugin on Azure version 3.0 supports service tags.

Azure Service tags simplify security for Azure virtual machines and Azure virtual networks because you can restrict network access to just the Azure services you want to use. A service tag represents a group of IP address prefixes for a particular Azure service. For example, a tag can represent all storage IP addresses.

The plugin makes a daily API call (at 5:00 am UTC) to retrieve all service tags from the Azure Portal, parses the payload to form IP-Service Mappings, and stores the mappings in the plugin database. The mappings are passed to configd, then on to Panorama. If the API call fails to return service information, the plugin forms the IP-Service mappings from the contents of **`service_tags_public.json`**. Plugin logs report the origin of the IP-Service mappings, the daily retrieval or the JSON file.

 ©2024 Palo Alto Networks, Inc.

The plugin also updates service tags for a new installation of the plugin, commit events, and monitoring definition addition or deletion.

A sample IP-Service mapping is shown below:

```
Service Name: AppServiceManagementazure.svc-tag.<service-name>
Example:
    azure.svc-tag.AppServiceManagement.WestUS2
Public IP CIDRs:
    13.166.40.0/26
    54.179.89.0/18
```

 ©2024 Palo Alto Networks, Inc.

Set up Active/Passive HA on Azure

You can configure a pair of VM-Series firewalls on Azure in an active/passive high availability (HA) configuration. For HA on Azure, you must deploy both firewall HA peers within the same Azure Resource Group and you must install the same version of the VM-Series Plugin on both HA peers.

- Set up Active/Passive HA on Azure (North-South & East-West Traffic)—If you have an internet-facing application deployed on your Azure infrastructure, and you need to secure north-south traffic, you require a floating IP address to secure traffic on failover. This floating IP address, which enables external connectivity, is always attached to the active peer. On failover, the process of detaching the IP address and reattaching it to the now active peer can take a few minutes.
- Set up Active/Passive HA on Azure (East-West Traffic Only)—If your application access and security requirements are contained within the Azure infrastructure and you need to secure east-west traffic only, you do not need a floating IP address. Instead, the HA implementation automatically reconfigures the UDRs in the Azure routing tables to provide a faster failover time.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

To enable HA on the VM-Series firewall on Azure, you must create an Azure Active Directory application and Service Principal that includes the permissions listed in the table below.

Azure HA Type	Permissions	Role Scope

Secondary IP Move HA	`"Microsoft.Authorization/*/read""Microsoft.Compute/virtualMachines/read""Microsoft.Network/networkInterfaces/*""Microsoft.Network/networkSecurityGroups/*""Microsoft.Network/virtualNetworks/join/action""Microsoft.Network/virtualNetworks/subnets/join/action"`
	The following permissions are required only if you have assigned a public IP address to any of your data interfaces. Standard SKU interface is recommended.
	`"Microsoft.Network/publicIPAddresses/join/action""Microsoft.Network/publicIPAddresses/`
		• Virtual network in which the VMs are deployed
		• Two VM-Series firewalls
		• NICs of both VM-Series firewalls
		• Network Security Group
		• Public IP addresses of the VM-Series firewalls

 ©2024 Palo Alto Networks, Inc.

Azure HA Type	Permissions	Role Scope
	`read""Microsoft.Network/publicIPAddresses/write"`	
UDR HA	`"Microsoft.Authorization/*/read""Microsoft.Compute/virtualMachines/read""Microsoft.Network/routeTables/*""Microsoft.Network/networkInterfaces/*"`	• Two VM-Series firewalls • NICs of both VM-Series firewalls • Route tables associated with UDR

Secondary IP Move and UDR	`"Microsoft.Authorization/*/read""Microsoft.Compute/virtualMachines/read""Microsoft.Network/networkInterfaces/*""Microsoft.Network/networkSecurityGroups/*""Microsoft.Network/routeTables/*""Microsoft.Network/virtualNetworks/join/action""Microsoft.Network/virtualNetworks/subnets/join/action"`
	The following permissions are required only if you have assigned a public IP address to any of your data interfaces. Standard SKU interface is recommended.
	`"Microsoft.Network/publicIPAddresses/join/action""Microsoft.Network/publicIPAddresses/read""Microsoft.Network/publicIPAddresses/write"`
		• Virtual network in which the VMs are deployed
		• Two VM-Series firewalls
		• NICs of both VM-Series firewalls
		• Network Security Group
		• Public IP addresses of the VM-Series firewalls
		• Route tables associated with UDR

Set up Active/Passive HA on Azure (North-South & East-West Traffic)

If you want to secure north-south traffic to your applications in your Azure infrastructure, use this workflow with floating IP addresses that can quickly move from one peer to the other. Because you cannot move the IP address associated with the primary interface of the firewall on Azure, you need to assign a secondary IP address that can function as a floating IP address. When the active firewall goes down, the floating IP address moves from the active to the passive firewall so that the passive firewall can seamlessly secure traffic as soon as it becomes the active peer. In addition to the floating IP address, the HA peers also need HA links—a control link (HA1) and a data link (HA2)—to synchronize data and maintain state information.

 ©2024 Palo Alto Networks, Inc.

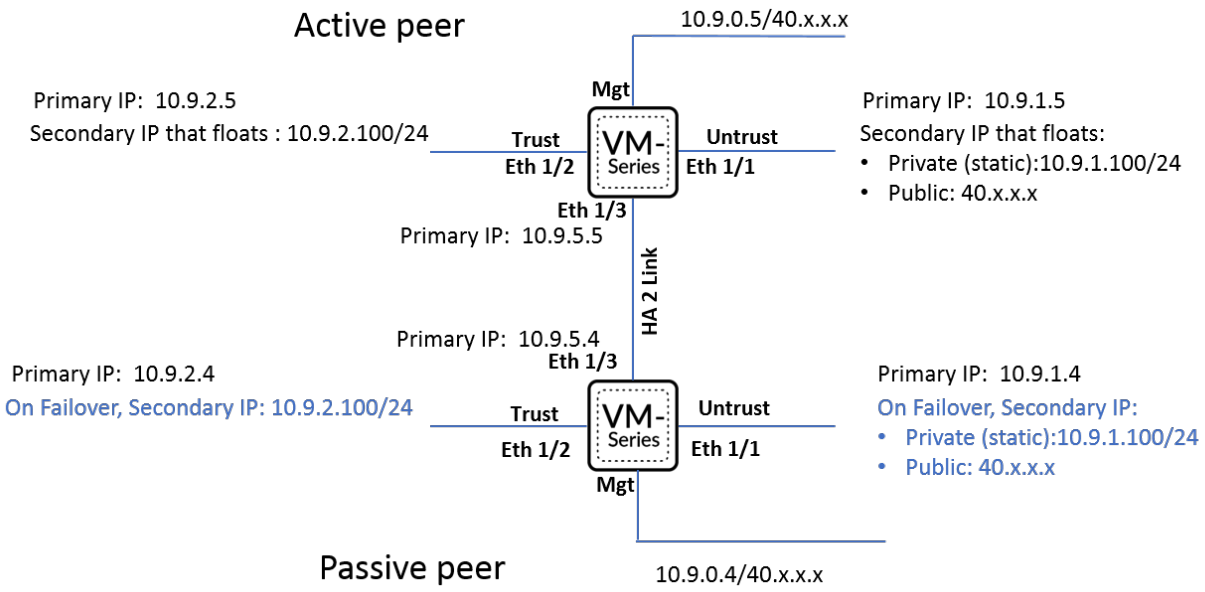

- Set up the Firewalls for Enabling HA
- Configure Active/Passive HA on the VM-Series Firewall on Azure

Set up the Firewalls for Enabling HA

Gather the following details for configuring HA on the VM-Series firewalls on Azure.

1. Set up the Active Directory application and a Service Principal to enable programmatic API access.
   - For the firewall to interact with the Azure APIs, you need to create an Azure Active Directory Service Principal. This Service Principle has the permissions required to authenticate to the Azure AD and access the resources within your subscription.To complete this set up, you must have permissions to register an application with your Azure AD tenant, and assign the application to a role in your subscription. If you don't have the necessary permissions, ask your Azure AD or subscription administrator to create a Service

 ©2024 Palo Alto Networks, Inc.

Principal. See the table above for the required permissions. Copy the following details for use later in this workflow:

- **Client ID**—The Application ID associated with the Active Directory (On the Azure portal, click **Home** > **Azure Active Directory** > **App registrations**, select your application and copy the ID).
- **Tenant ID**—The Directory ID associated with the Active Directory (On the Azure portal, click **Home** > **Azure Active Directory** > **Properties** > **Directory ID**, select the application and copy the ID).

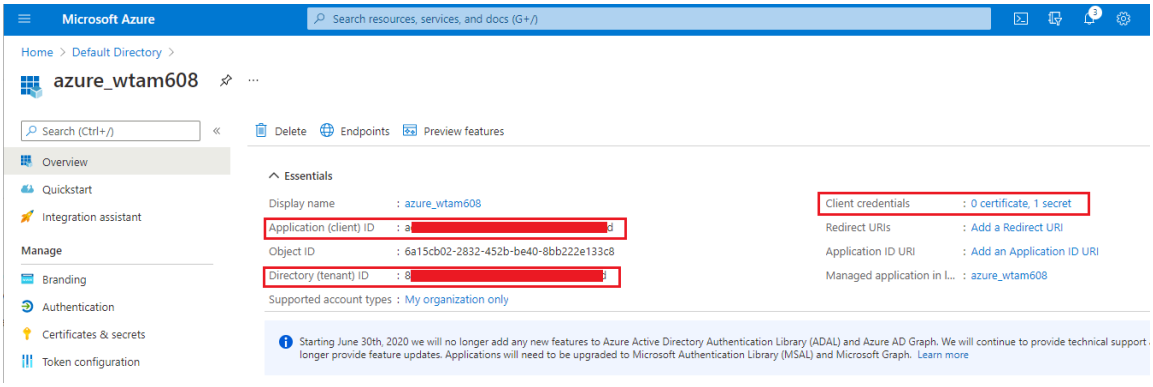

- **Azure Subscription ID**—The Azure subscription in which you have deployed the firewalls. You must login to your Azure portal to get this subscription ID.
- **Resource Group Name**— The resource group name in which you have deployed the firewalls that you want to configure as HA peers. Both firewalls must be in the same resource group.

 ©2024 Palo Alto Networks, Inc.

- **Secret Key**—The authentication key associated with the Active Directory application (On the Azure portal, click **Home** > **Azure Active Directory** > **Certificates & secrets**, copy the **Value** under **Client secrets**. If you do not have a Secret Key, create one first, then

 ©2024 Palo Alto Networks, Inc.

copy the value). To log in as the application, you must provide both the key value and the Application ID.

 ©2024 Palo Alto Networks, Inc.

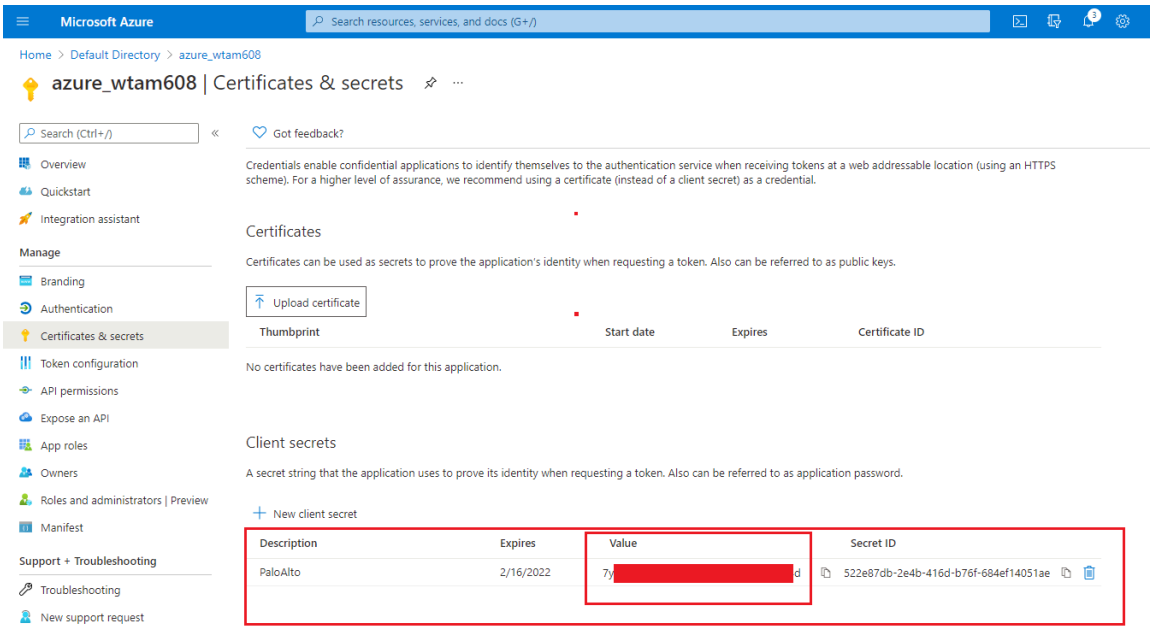

 ©2024 Palo Alto Networks, Inc.

- Know where to get the templates you need to deploy the VM-Series firewalls within the same Azure Resource Group.

  For an HA configuration, both HA peers must belong to the same Azure Resource Group. If you deploy the first instance of the firewall from the Azure Marketplace, and must use your custom ARM template for deploying the second instance of the firewall into the existing Resource Group. The reason you need a custom template or the Palo Alto Networks sample template is because Azure does not support the ability to deploy the firewall in to an Resource Group that is not empty.

  Copy the deployment information for the first firewall instance. For example:

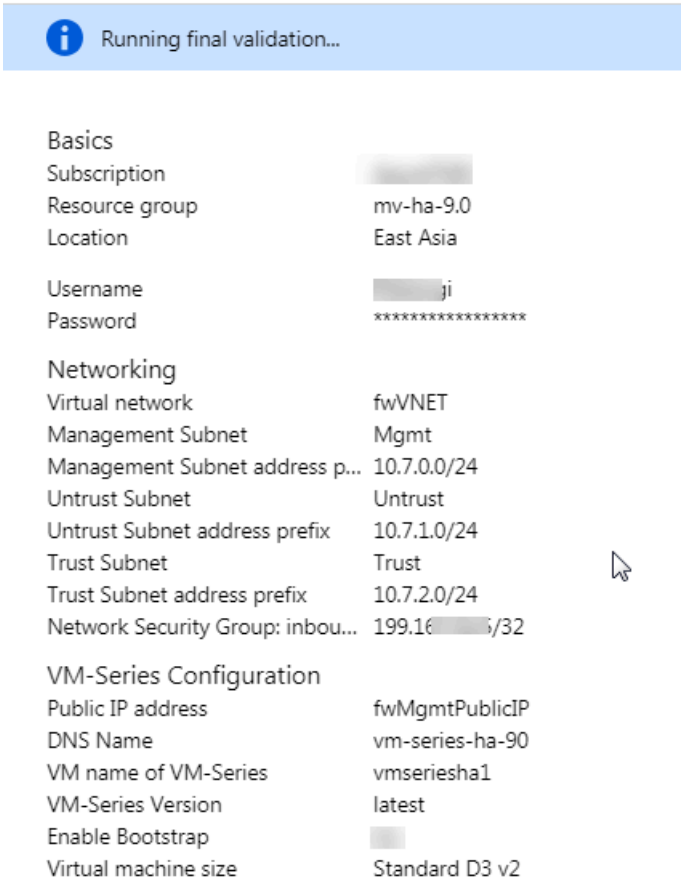

- Match the **VM Name of VM-Series** firewall as shown in the screenshot above with the **Hostname** on the firewall web interface. You must add the same name on **Device** > **Setup** > **Management**, because the hostname of the firewall is used to trigger failover.

 ©2024 Palo Alto Networks, Inc.

- Plan the network interface configuration on the VM-Series firewalls on Azure.

  To set up HA, you must deploy both HA peers within the same Azure Resource Group and both firewalls must have the same number of network interfaces. A minimum of four network interfaces is required on each HA peer:

  - **Management interface (eth0)**—Private and public IP address associated with the primary interface. The public IP address enables access to the firewall web interface and SSH access.

    You can use the private IP interface on the management interface as the HA1 peer IP address for the control link communication between the active/passive HA peers. If you want a dedicated HA1 interface, you must attach an additional network interface on each firewall, and this means that you need five interfaces on each firewall.
  - **Untrust interface (eth1/1)**—Primary private IP address with /32 netmask, and secondary IP configuration with both a private IP address (any netmask) and a public IP address.

    On failover, when the passive peer transitions to the active state, the public IP address associated with the secondary IP configuration is detached from the previously active peer and attached to the now active HA peer.
  - **Trust interface (eth1/2)**—Primary and secondary private IP addresses. On failover, when the passive peer transitions to the active state, the secondary private IP address is detached from the previously active peer and is attached to the now active HA peer.
  - **HA2 (eth 1/3)**—Primary private IP address. The HA2 interface is the data link that the HA peers use for synchronizing sessions, forwarding tables, IPSec security associations and ARP tables.

Interface	Active firewall peer	Passive firewall peer	Description
Trust	Secondary IP address	—	The trust interface of the active peer requires a secondary IP configuration that can float to the other peer on failover. This secondary IP configuration on the trust interface must be a private IP address with the netmask of the servers that it secures. On failover, the VM-Series plugin calls the Azure API to detach this secondary private IP address from the active peer and attach it to the passive peer. Attaching this IP address to the now active peer ensures that the firewall can receive traffic on the floating IP on the untrust interface and send it through to the floating IP on the trust interface and on to the workloads.
Untrust	Secondary IP address	—	The untrust interface of the firewall requires a secondary IP configuration that includes a static private IP address with a netmask for the untrust subnet, and a public IP address for accessing the back-end servers or workloads over the internet. On failover, the VM-Series plugin

 ©2024 Palo Alto Networks, Inc.

Interface	Active firewall peer	Passive firewall peer	Description
			calls the Azure API to detach the secondary IP configuration from the active peer and attach it to the passive peer before it transitions to the active state. This process of floating the secondary IP configuration, enables the now active firewall to continue processing inbound traffic that is destined to the workloads.
HA2	Add a NIC to the firewall from the Azure management console.	Add a NIC to the firewall from the Azure management console.	On the active and passive peers, add a dedicated HA2 link to enable session synchronization. The default interface for HA1 is the management interface, and you can opt to use the management interface instead of adding an additional interface to the firewall. For enabling data flow over the HA2 link, you need to add an additional network interface on the Azure portal and configure the interface for HA2 on the firewall.

Configure Active/Passive HA on the VM-Series Firewall on Azure

In this workflow, you deploy the first instance of the VM-Series firewall using the VM-Series firewall solution template in the Azure marketplace, and the second instance of the firewall using the sample GitHub template.

The authentication key (client secret) associated with the Active Directory application required for setting up the VM-Series firewall in an HA configuration, is encrypted with VM-Series plugin version 1.0.4 on the firewall and on Panorama. Because the key is encrypted in VM-Series plugin version 1.0.4, you must install the same version of the plugin on Panorama and the managed VM-Series firewalls in order to centrally manage the firewalls from Panorama.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Deploy the VM-Series firewall using a solution template and set up the network interfaces for HA.

1. Add a secondary IP configuration to the untrust interface of the firewall.

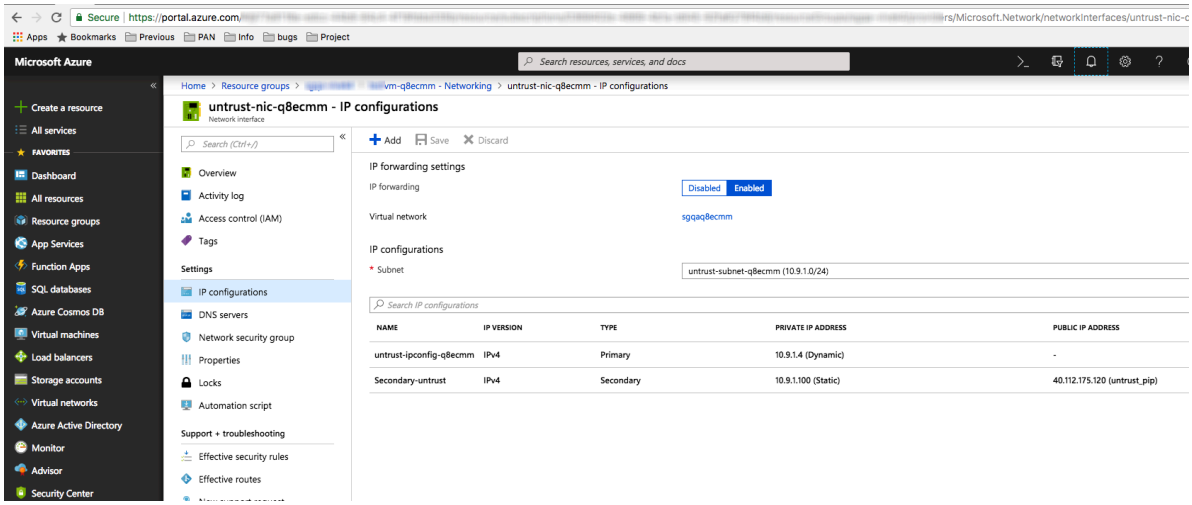

You must attach the secondary IP configuration—with a private IP address (any netmask) and a public IP address—to the firewall that will be designated as the active peer. The secondary IP configuration always stays with the active HA peer, and moves from one peer to the another when a failover occurs.

In this workflow, this firewall will be designated as the active peer. The active HA peer has a lower numerical value for device priority that you configure as a part of the HA

 ©2024 Palo Alto Networks, Inc.

configuration on the firewall, and this value indicates a preference for which firewall assumes the role of the active peer.

2. Add a secondary IP configuration to the trust interface of the firewall.

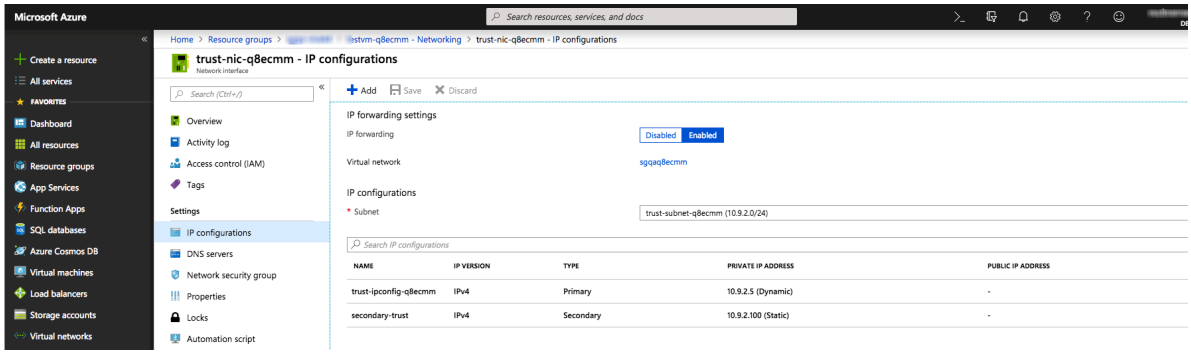

The secondary IP configuration for the trust interface requires a static private IP address only. This IP address moves from the active firewall to the passive firewall on failover so

 ©2024 Palo Alto Networks, Inc.

that traffic flows through from the untrust to the trust interface and to the destination subnets that the firewall secures.

3. Attach a network interface for the HA2 communication between the firewall HA peers.
   1. Add a subnet within the virtual network.
   2. Create and attach a network interface to the firewall.
4. Set up your route table on Azure.

   Your next hop should point to the floating IP address as shown here:

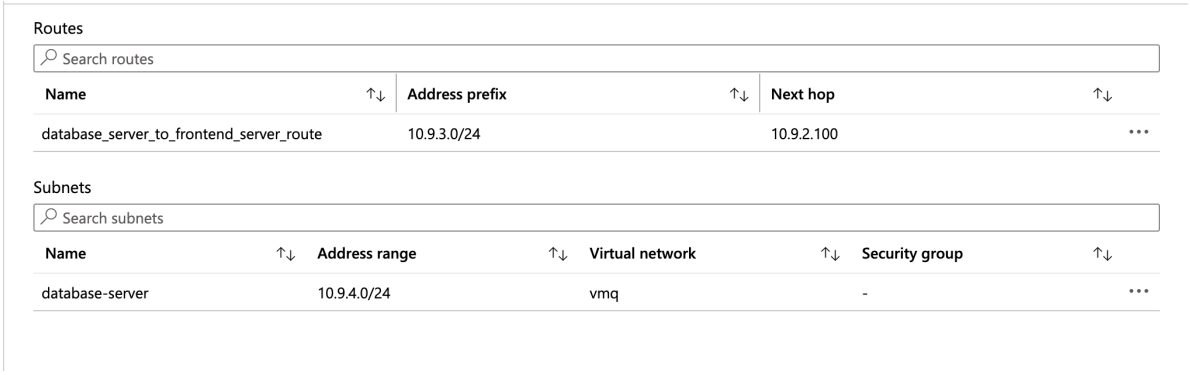

 ©2024 Palo Alto Networks, Inc.

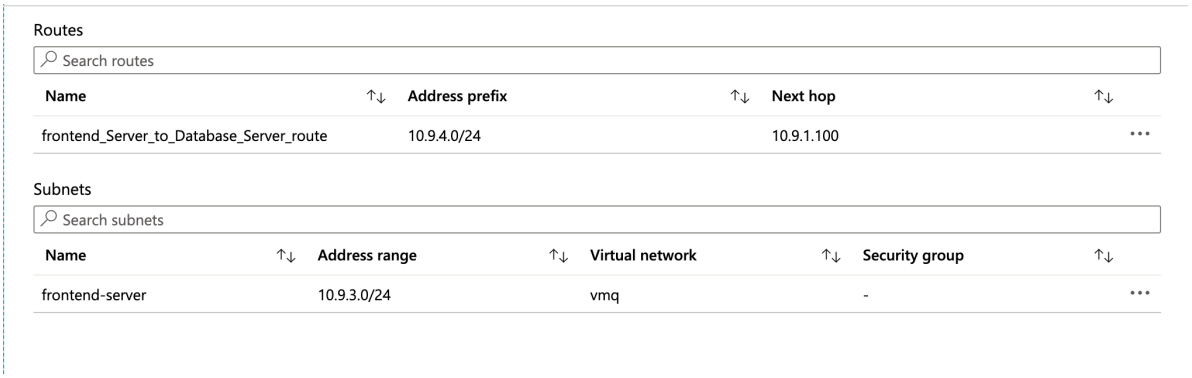
Routes

Search routes

Name	Address prefix	Next hop
frontend_Server_to_Database_Server_route	10.9.4.0/24	10.9.1.100

Subnets

Search subnets

Name	Address range	Virtual network	Security group
frontend-server	10.9.3.0/24	vmq	-

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure the interfaces on the firewall.

Complete these steps on the active HA peer, before you deploy and set up the passive HA peer.

1. Log in to the firewall web interface.
2. Configure ethernet 1/1 as the untrust interface and ethernet 1/2 as the trust interface.

   Select **Network** > **Interfaces** and configure as follows:

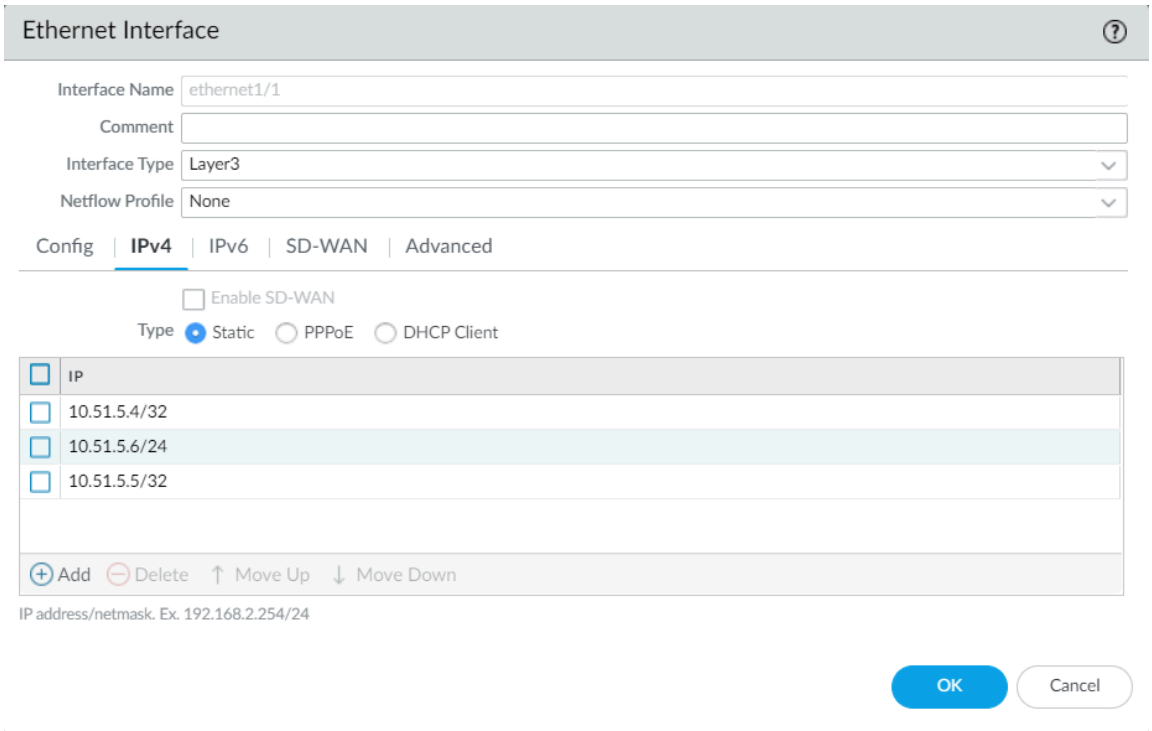

 ©2024 Palo Alto Networks, Inc.

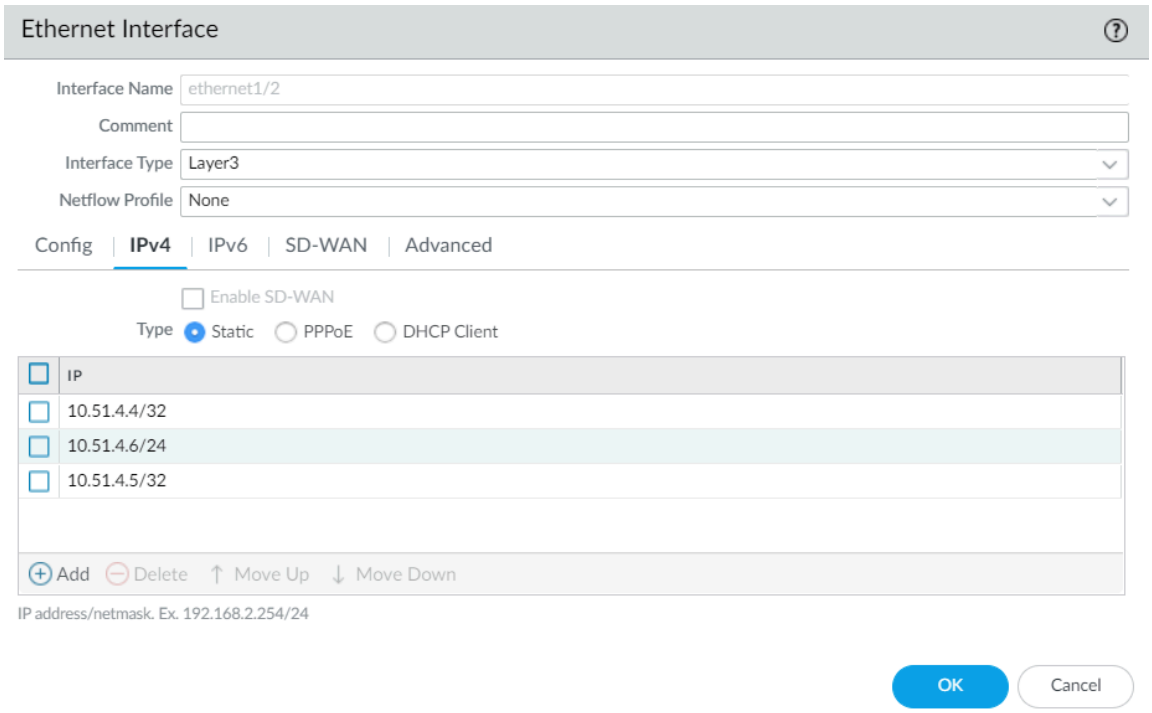

3. Configure ethernet 1/3 as the HA interface.

   To set up the HA2 link, select the interface and set **Interface Type** to **HA**. Set link speed and duplex to auto.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure the VM-Series plugin to authenticate to the Azure resource group in which you have deployed the firewall.

Set up the Azure HA configuration on the VM-Series plugin.

To encrypt the client secret, use the VM-Series plugin version 1.0.4 or later. If using Panorama to manage your firewalls, you must install the VM-Series plugin version 1.0.4 or later.

1. Select **Device** > **VM-Series** to enable programmatic access between the firewall plugin and the Azure resources.

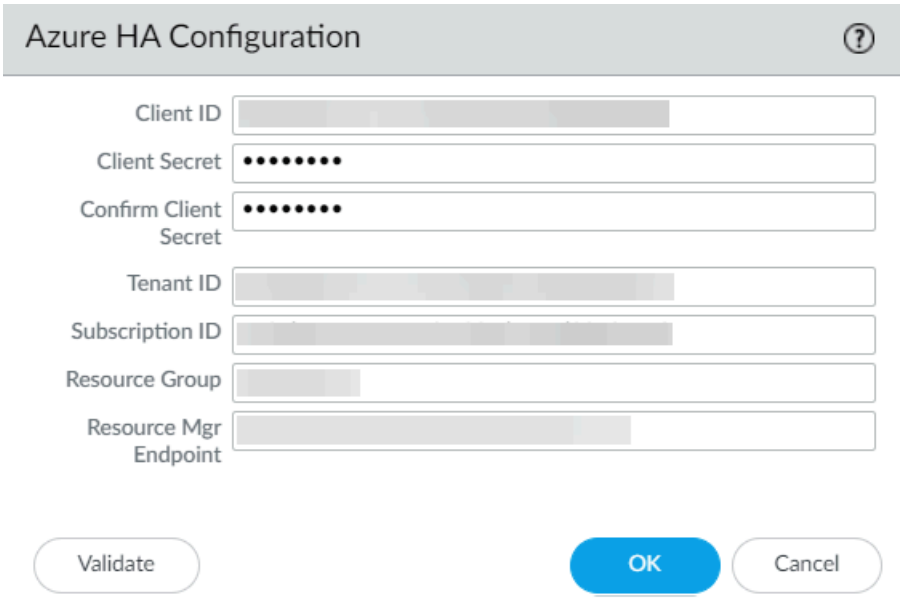

2. Enter the **Client ID**. The client ID is the Application ID associated with your Azure Active Directory application.
3. Enter the **Client Secret** and re-enter it to confirm.
4. Enter the **Tenant ID**. The tenant ID is the Directory ID you saved when you set up the Active Directory application.
5. Enter the **Subscription ID** for the Azure subscription you want to monitor.
6. Enter the **Resource Group** name.
7. (For Azure Stack deployments only) Enter the **Resource Mgr Endpoint** URL. This field is mandatory ONLY for Azure Stack deployments. Do not enter a value for this field if you are using a regular Azure Cloud deployment; HA failover will not succeed if you specify the **Resource Mgr Endpoint** URL for a regular Azure Cloud deployment.

   This field is available in VM-Series plugin 2.1.2 and later.

8. Click **Validate** to verify that the keys and IDs you entered are valid, and that VM-Series plugin can successfully communicate with the Azure resources using the API.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enable HA.

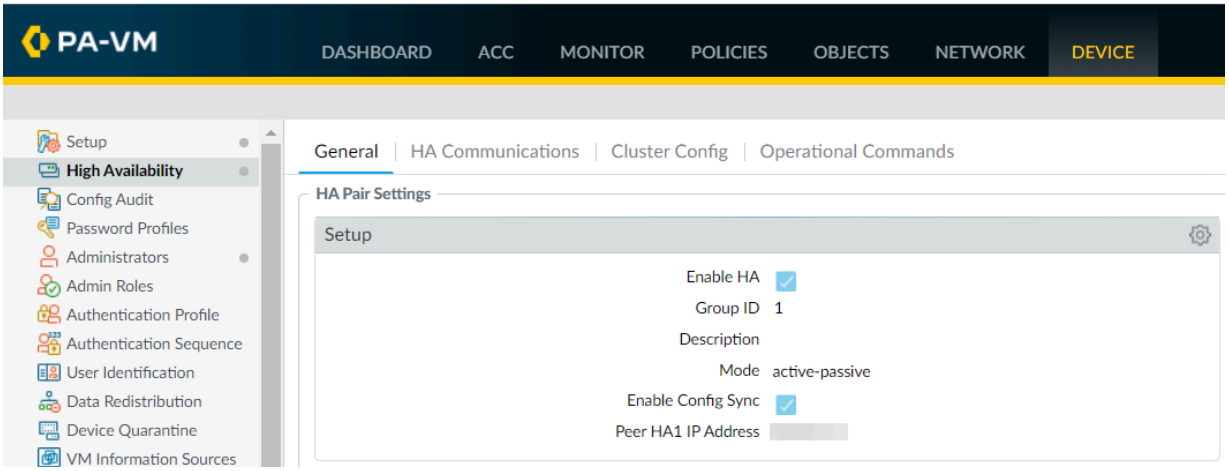

1. Select **Device** > **Setup** > **HA**.
2. Enter **Peer HA1 IP address** as the private IP address of the passive peer.
3. (Optional) Edit the Control Link (HA1). If you do not plan to use the management interface for the control link and have added an additional interface (for example ethernet 1/4), edit this section to select the interface to use for HA1 communication.
4. Edit the Data Link (HA2) to use **Port** ethernet 1/3 and add the IP address of this peer and the **Gateway** IP address for the subnet.

STEP 5 | **Commit** the changes.

STEP 6 | Set up the passive HA peer within the same Azure Resource Group.

1. Deploy the second instance of the firewall.
   - Download the custom template and parameters file from GitHub.
   - Log in to the Azure Portal.
   - Search for **custom template** and select **Deploy from a custom template**.
   - Select **Build your own template in the editor** > **Load file**.
   - Select the **azuredeploy.json** that you downloaded earlier, and **Save**.
   - Complete the inputs, agree to the terms and **Purchase**.

     Make sure to match the following inputs to that of the firewall instance you have already deployed— Azure subscription, name of the Resource Group, location of the Resource Group, name of the existing VNet into which you want to deploy the

 ©2024 Palo Alto Networks, Inc.

firewall, VNet CIDR, Subnet names, Subnet CIDRs, and start the IP address for the management, trust and untrust subnets.

2. Repeat Step 1and Step 2to set up the interfaces and configure the firewall as the passive HA peer.
3. Skip Step 3 and complete Enable HA (Step 5). In Step 4 modify the IP addresses as appropriate for this passive HA peer.

STEP 7 | After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls, and view the High Availability widget.
2. On the active firewall, click the **Sync to peer** link.
3. Confirm that the firewalls are paired and synced, as shown as follows:
   - On the passive firewall: the state of the local firewall should display **passive** and the **Running Config** should show as **synchronized**.
   - On the active firewall: The state of the local firewall should display **active** and the **Running Config** should show as **synchronized**.
4. On the passive peer, verify that the VM-Series plugin configuration is now synced.

   Select **Device** > **VM-Series** and validate that you can view the Azure HA configuration that you had omitted configuring on the passive peer.

Set up Active/Passive HA on Azure (East-West Traffic Only)

If your resources are all deployed within the Azure infrastructure and you do not need to enforce security for north south traffic to the Azure VNet, you can deploy a pair of VM-Series firewalls in an active/passive high availability (HA) configuration without floating IP addresses. The HA peers will still need HA links—a control link (HA1) and a data link (HA2)—to synchronize data and maintain state information.

You must have the VM-Series Plugin version 1.0.9 or later, and you must deploy both firewall HA peers within the same Azure Resource Group.

 ©2024 Palo Alto Networks, Inc.

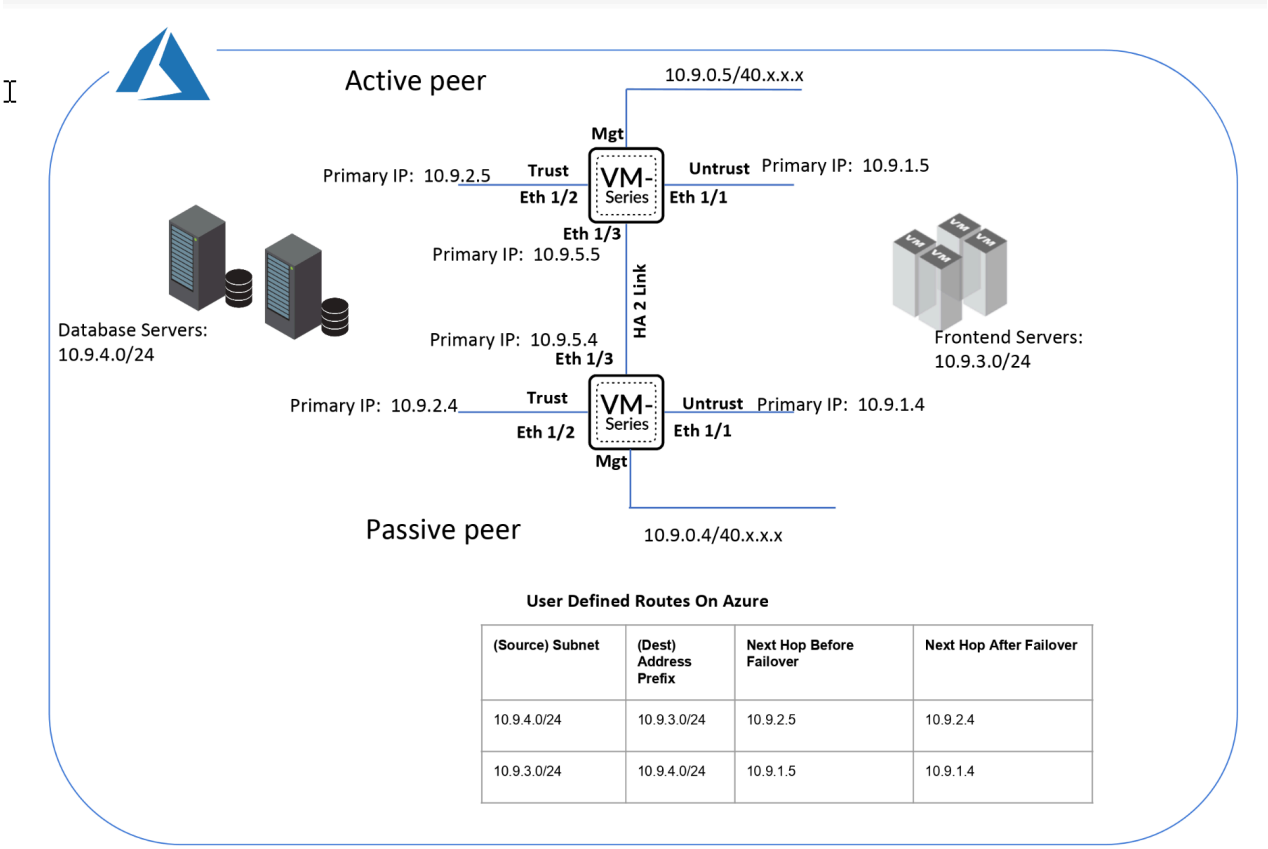

User Defined Routes On Azure

(Source) Subnet	(Dest) Address Prefix	Next Hop Before Failover	Next Hop After Failover
10.9.4.0/24	10.9.3.0/24	10.9.2.5	10.9.2.4
10.9.3.0/24	10.9.4.0/24	10.9.1.5	10.9.1.4

- Set up the Firewalls for Enabling HA
- Configure Active/Passive HA on the VM-Series Firewall on Azure

Set up the Firewalls for Enabling HA

Gather the following details for configuring HA on the VM-Series firewalls on Azure.

- Set up the Active Directory application and a Service Principal to enable programmatic API access.
  - For the firewall to interact with the Azure APIs, you need to create an Azure Active Directory Service Principal. This Service Principle has the permissions required to authenticate to the Azure AD and access the resources within your subscription.To complete this set up, you must have permissions to register an application with your Azure AD tenant, and assign the application to a role in your subscription. If you don't have the necessary permissions, ask your Azure AD or subscription administrator to create a Service

 ©2024 Palo Alto Networks, Inc.

Principal. See the table above for the required permissions. Copy the following details for use later in this workflow:

- **Client ID**—The Application ID associated with the Active Directory (On the Azure portal, click **Home** > **Azure Active Directory** > **App registrations**, select your application and copy the ID).
- **Tenant ID**—The Directory ID associated with the Active Directory (On the Azure portal, click **Home** > **Azure Active Directory** > **Properties** > **Directory ID**, select the application and copy the ID).

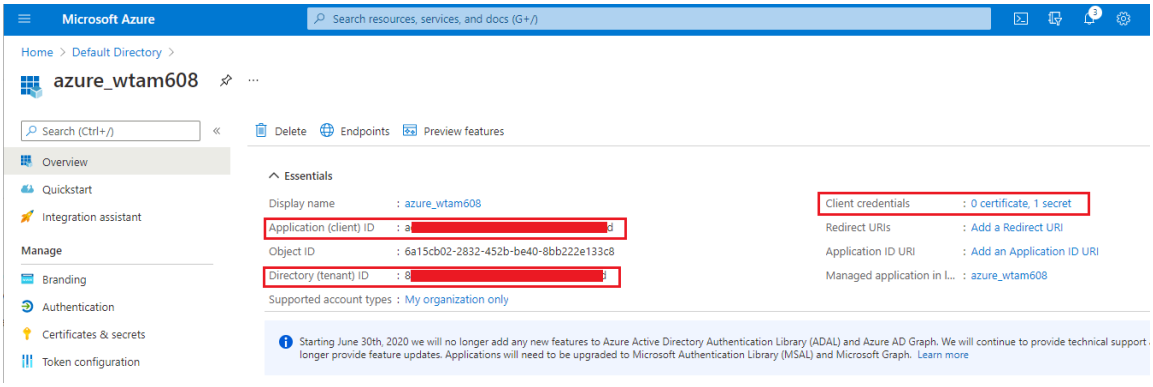

- **Azure Subscription ID**—The Azure subscription in which you have deployed the firewalls. You must login to your Azure portal to get this subscription ID.
- **Resource Group Name**— The resource group name in which you have deployed the firewalls that you want to configure as HA peers. Both firewalls must be in the same resource group.

 ©2024 Palo Alto Networks, Inc.

- **Secret Key**—The authentication key associated with the Active Directory application (On the Azure portal, click **Home** > **Azure Active Directory** > **Certificates & secrets**, copy the **Value** under **Client secrets**. If you do not have a Secret Key, create one first, then

 ©2024 Palo Alto Networks, Inc.

copy the value). To log in as the application, you must provide both the key value and the Application ID.

 ©2024 Palo Alto Networks, Inc.

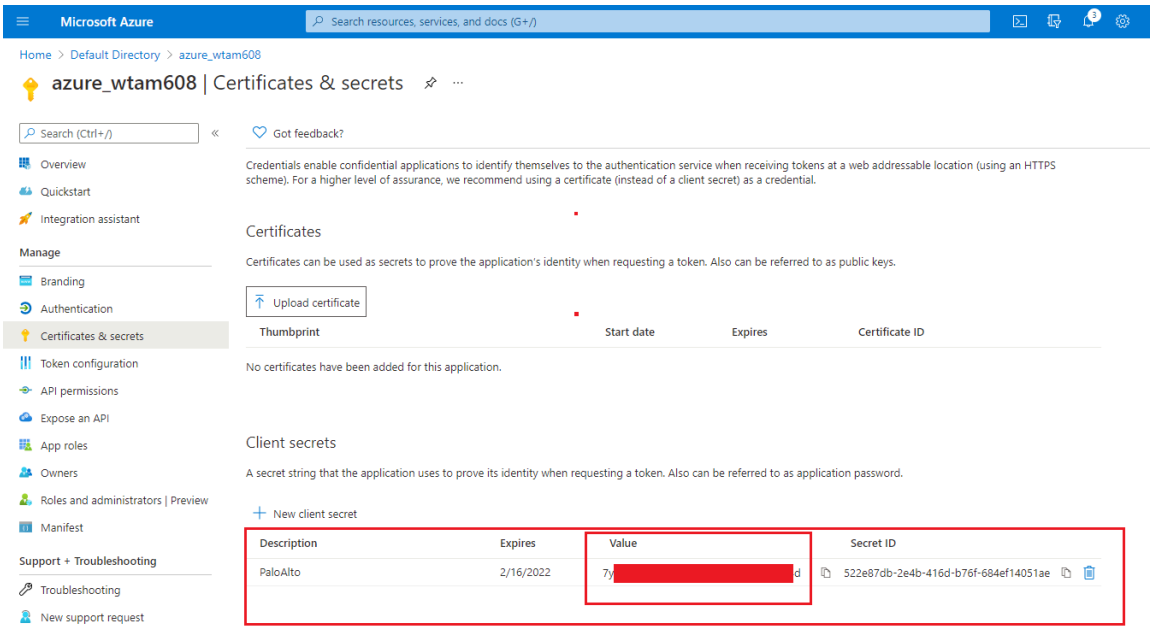

 ©2024 Palo Alto Networks, Inc.

- Know where to get the templates you need to deploy the VM-Series firewalls within the same Azure Resource Group.

  For an HA configuration, both HA peers must belong to the same Azure Resource Group. If you deploy the first instance of the firewall from the Azure Marketplace, and must use your custom ARM template or the Palo Alto Networks sample GitHub template for deploying the second instance of the firewall into the existing Resource Group. The reason you need a custom template or the Palo Alto Networks sample template is because Azure does not support the ability to deploy the firewall in to an Resource Group that is not empty.

  Copy the deployment information for the first firewall instance. For example:

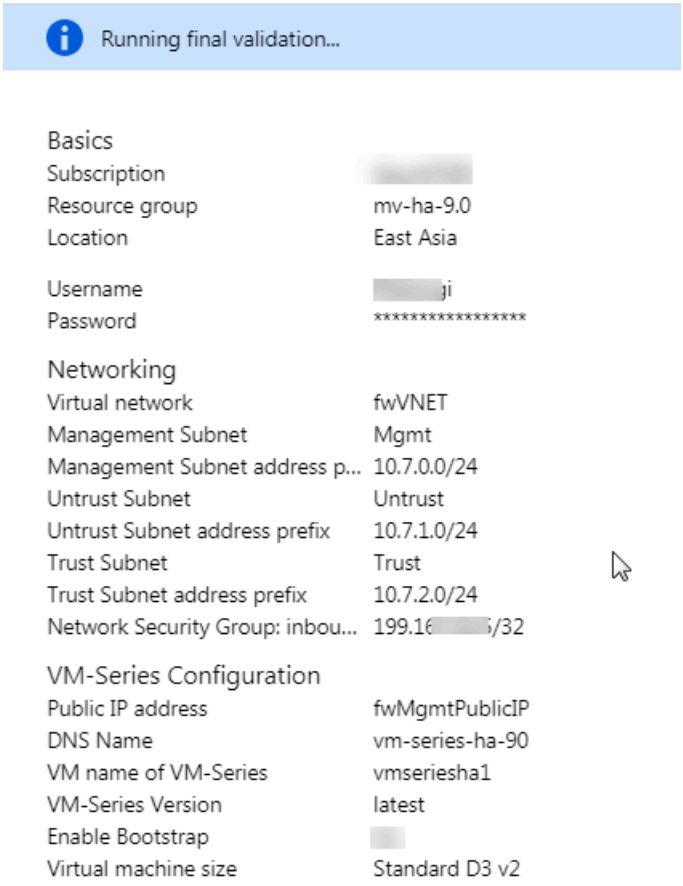

- Match the **VM Name of VM-Series** firewall as shown in the screenshot above with the **Hostname** on the firewall web interface. You must add the same name on **Device** > **Setup** > **Management**, because the hostname of the firewall is used to trigger failover.

 ©2024 Palo Alto Networks, Inc.

- Plan the network interface configuration on the VM-Series firewalls on Azure.

  To set up HA, you must deploy both HA peers within the same Azure Resource Group and both firewalls must have the same number of network interfaces. A minimum of four network interfaces is required on each HA peer:

  - **Management interface (eth0)**—Private and public IP address associated with the primary interface. The public IP address enables access to the firewall web interface and SSH access.

    You can use the private IP interface on the management interface as the HA1 peer IP address for the control link communication between the active/passive HA peers. If you want a dedicated HA1 interface, you must attach an additional network interface on each firewall, and this means that you need five interfaces on each firewall.
  - **Untrust interface (eth1/1)**—Primary private IP address with /32 netmask.

    On failover, when the passive peer transitions to the active state, the VM-Series plugin automatically sends traffic to the primary private IP address of the passive peer. The Azure UDRs enable the traffic flow.
  - **Trust interface (eth1/2)**—Primary private IP address. On failover, when the passive peer transitions to the active state, the VM-Series plugin automatically sends traffic to the primary private IP address of the passive peer.
  - **HA2 (eth 1/3)**—Primary private IP address. The HA2 interface is the data link that the HA peers use for synchronizing sessions, forwarding tables, IPSec security associations and ARP tables.

Interface	Active firewall peer	Passive firewall peer	Description
HA2	Add a NIC to the firewall from the Azure management console.	Add a NIC to the firewall from the Azure management console.	On the active and passive peers, add a dedicated HA2 link to enable session synchronization. The default interface for HA1 is the management interface, and you can opt to use the management interface instead of adding an additional interface to the firewall. For enabling data flow over the HA2 link, you need to add an additional network interface on the Azure portal and configure the interface for HA2 on the firewall.

Configure Active/Passive HA on the VM-Series Firewall on Azure

In this workflow, you deploy the first instance of the VM-Series firewall using the VM-Series firewall solution template in the Azure marketplace, and the second instance of the firewall using the sample GitHub template.

 ©2024 Palo Alto Networks, Inc.

The authentication key (client secret) associated with the Active Directory application required for setting up the VM-Series firewall in an HA configuration, is encrypted with VM-Series plugin version 1.0.9 on the firewall and on Panorama. Because the key is encrypted in VM-Series plugin version 1.0.9, you must install the same version of the plugin on Panorama and the managed VM-Series firewalls in order to centrally manage the firewalls from Panorama.

STEP 1 | Deploy the VM-Series firewall using a solution template and set up the network interfaces for HA.

For securing east west traffic within an Azure VNet, you only need a primary IP address for the trust and untrust firewall interfaces. When a failover occurs, the UDR changes and the route points to the primary IP address of the peer that transitions to the active state.

1. Add a Primary IP configuration to the trust interface of the active firewall peer.

   In this workflow, this firewall will be designated as the active peer. The active HA peer has a lower numerical value for device priority that you configure as a part of the HA

 ©2024 Palo Alto Networks, Inc.

configuration on the firewall, and this value indicates a preference for which firewall assumes the role of the active peer.

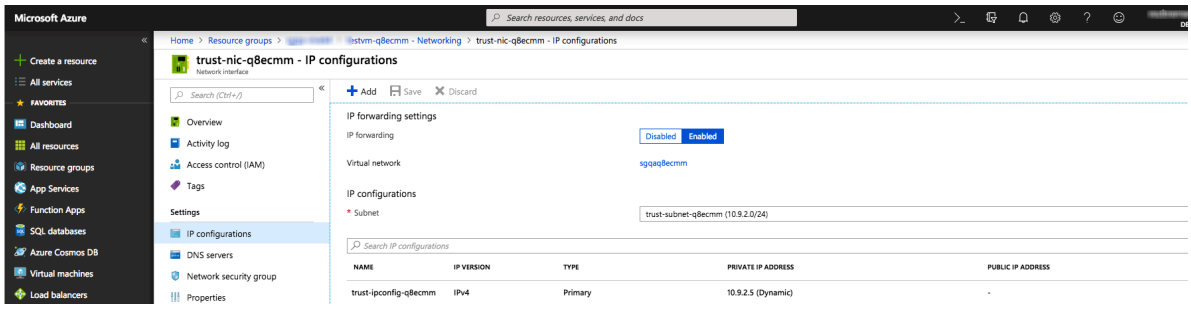

2. Add a Primary IP configuration to the untrust interface of the active firewall peer.

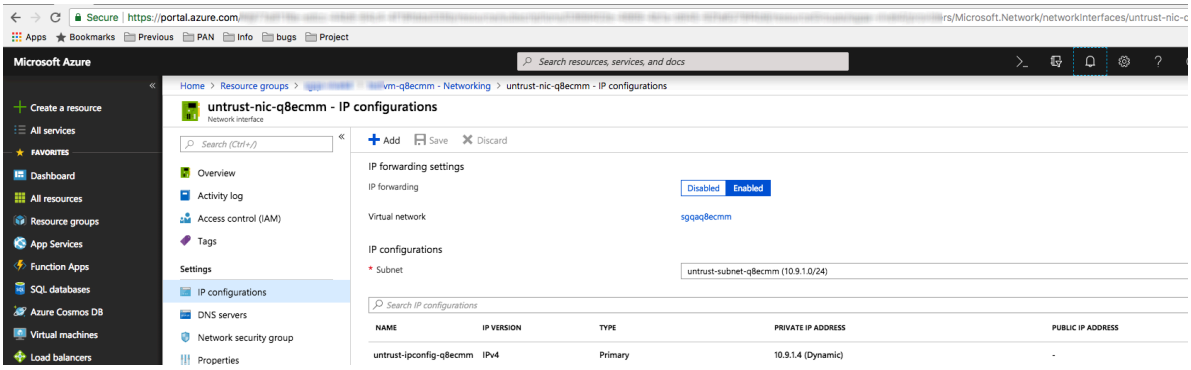

3. Attach a network interface for the HA2 communication between the firewall HA peers.

 ©2024 Palo Alto Networks, Inc.

1. Add a subnet within the virtual network.
2. Create and attach a network interface to the firewall.

4. Set up your route table on Azure.

   Create a route to the Next hop of Primary IP address of the trust and untrust interfaces of the active firewall peer.

Example with Frontend Server to Database Server Route :

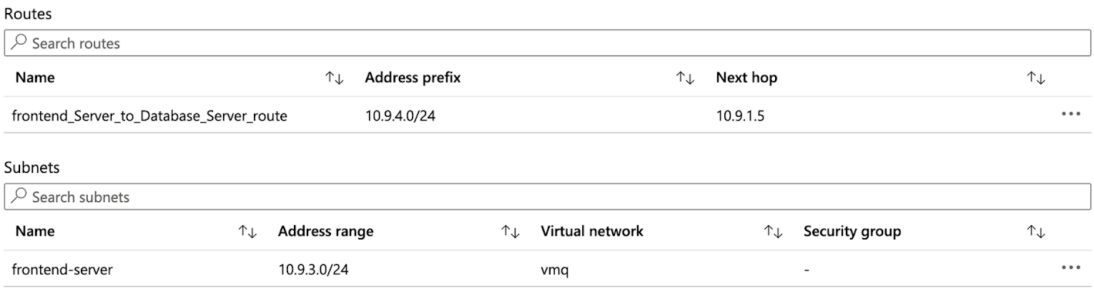

Example with Database Server to Frontend Server route:

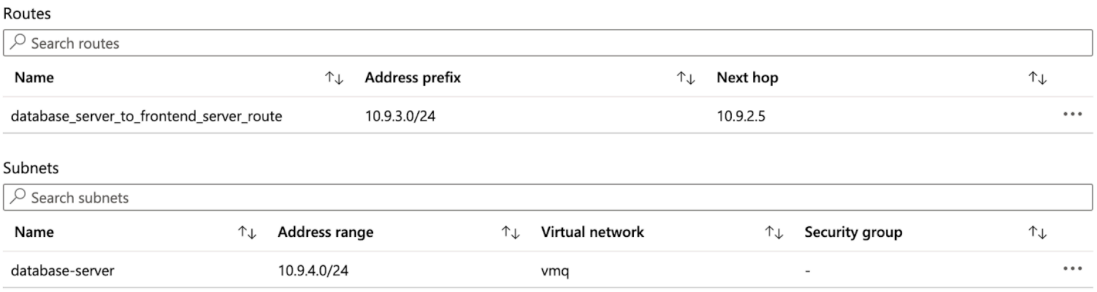

After failover, the next hop for the Database server to Frontend server route will change from 10.9.2.5 to 10.9.2.4. Similarly, the next hop of Frontend server to Database server route will change from 10.9.1.5 to 10.9.1.4.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure the interfaces on the firewall.

Complete these steps on the active HA peer, before you deploy and set up the passive HA peer.

1. Log in to the firewall web interface.
2. Configure ethernet 1/1 as the untrust interface and ethernet 1/2 as the untrust interface.

   Select **Network** > **Interfaces** and configure as follows:

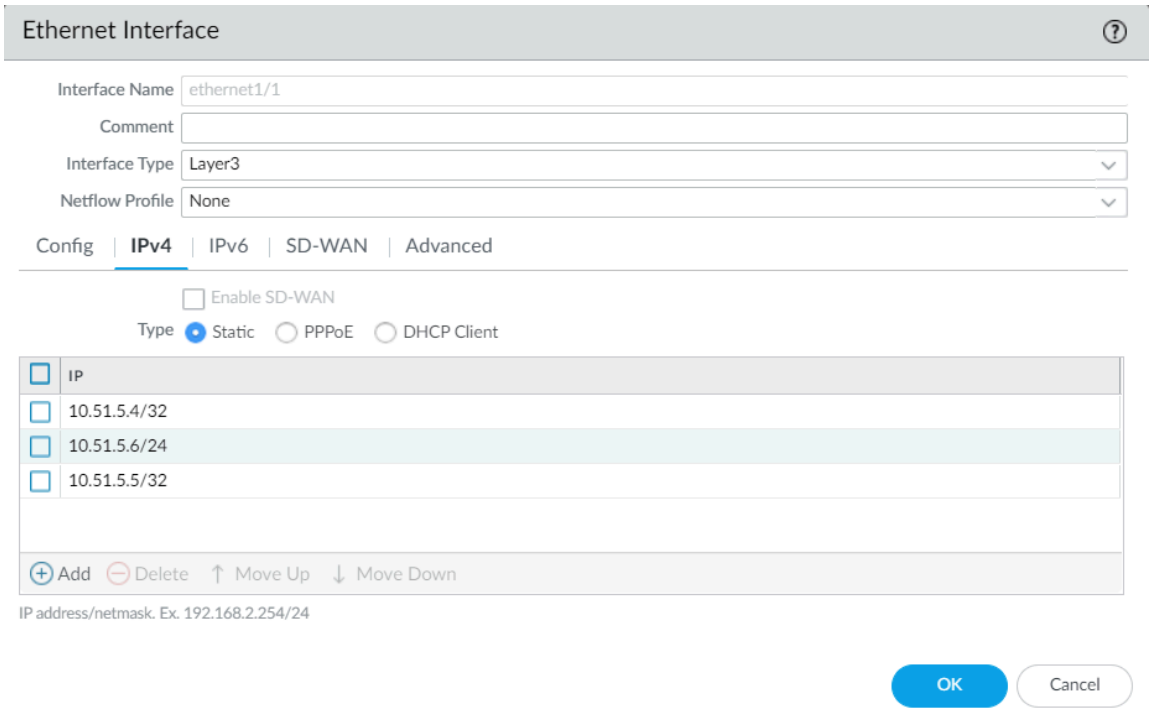

 ©2024 Palo Alto Networks, Inc.

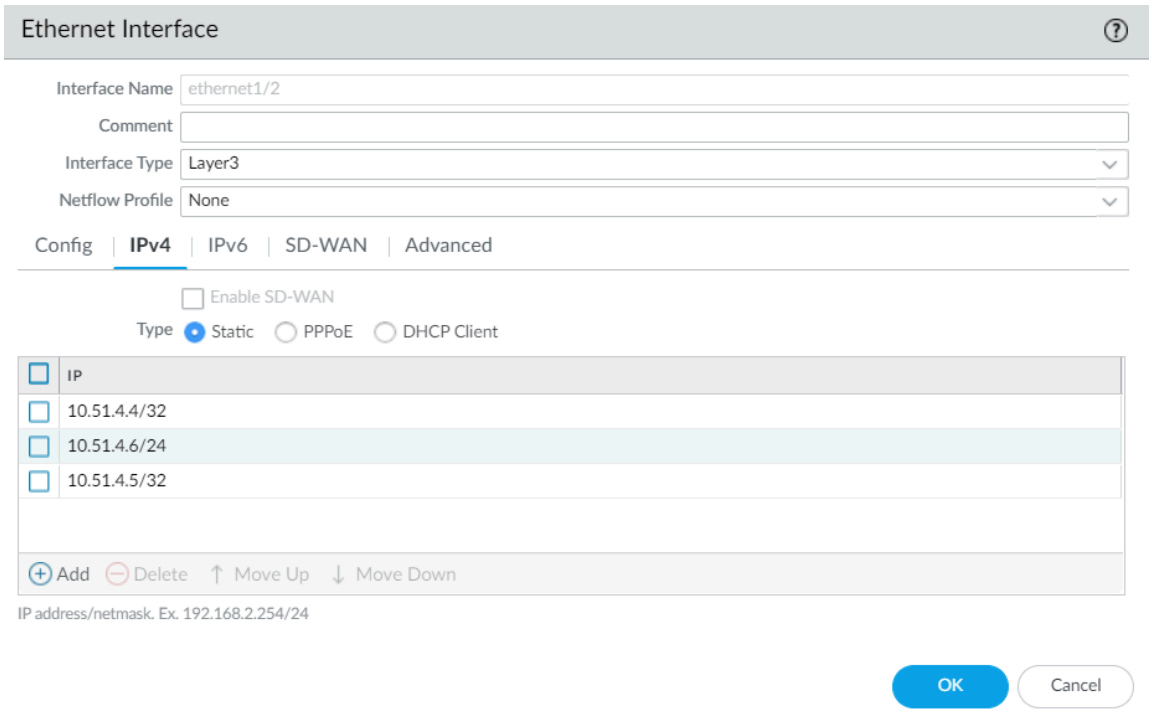

3. Configure ethernet 1/3 as the HA interface.

   To set up the HA2 link, select the interface and set **Interface Type** to **HA**. Set link speed and duplex to auto.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure the VM-Series plugin to authenticate to the Azure resource group in which you have deployed the firewall.

Set up the Azure HA configuration on the VM-Series plugin.

To encrypt the client secret, use the VM-Series plugin version 1.0.4 or later. If using Panorama to manage your firewalls, you must install the VM-Series plugin version 1.0.4 or later.

1. Select **Device** > **VM-Series** to enable programmatic access between the firewall plugin and the Azure resources.

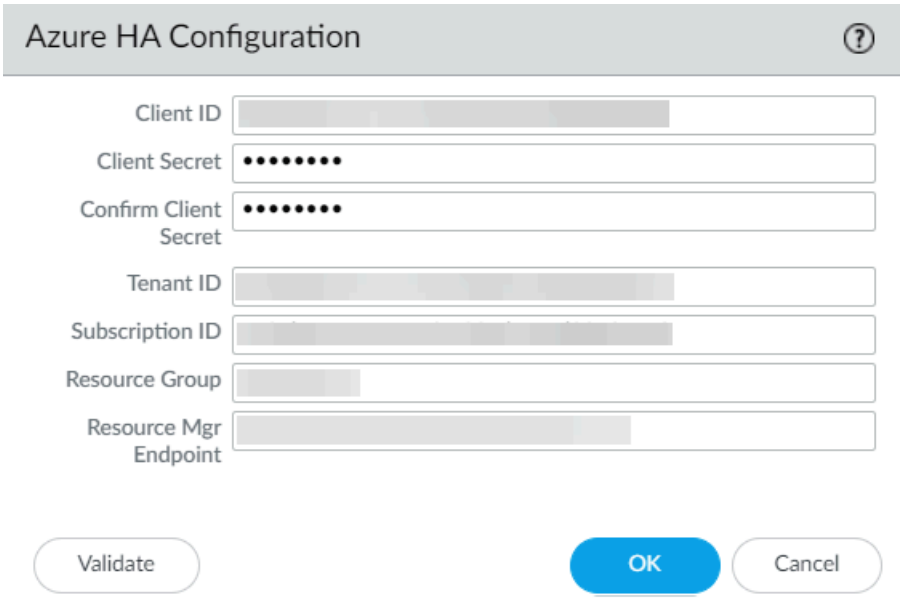

2. Enter the **Client ID**. The client ID is the Application ID associated with your Azure Active Directory application.
3. Enter the **Client Secret** and re-enter it to confirm.
4. Enter the **Tenant ID**. The tenant ID is the Directory ID you saved when you set up the Active Directory application.
5. Enter the **Subscription ID** for the Azure subscription you want to monitor.
6. Enter the **Resource Group** name.
7. (For Azure Stack deployments only) Enter the **Resource Mgr Endpoint** URL.

   This field is available in VM-Series plugin 2.1.2 and later.

8. Click **Validate** to verify that the keys and IDs you entered are valid, and that VM-Series plugin can successfully communicate with the Azure resources using the API.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enable HA.

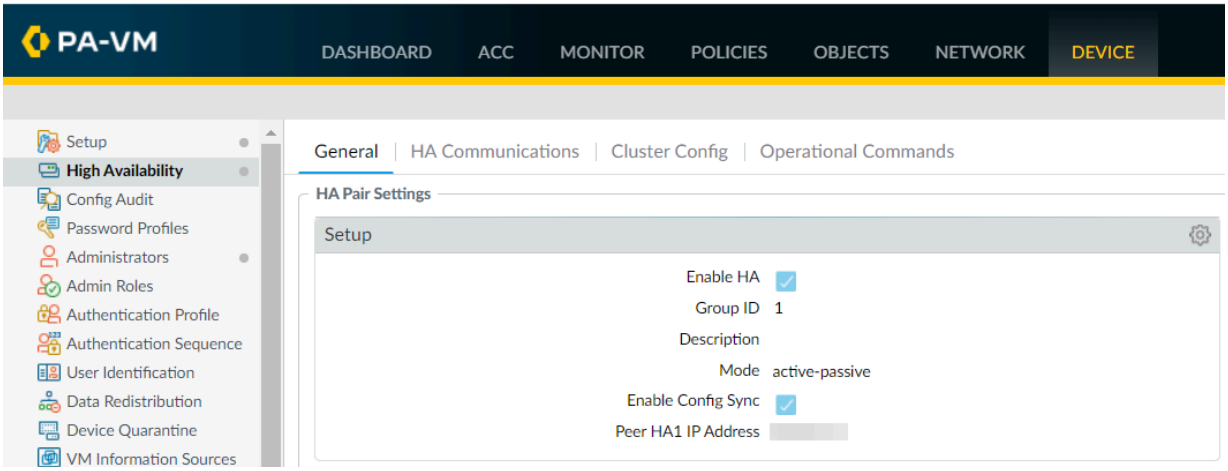

1. Select **Device** > **Setup** > **HA**.
2. Enter **Peer HA1 IP address** as the private IP address of the passive peer.
3. (Optional) Edit the Control Link (HA1). If you do not plan to use the management interface for the control link and have added an additional interface (for example ethernet 1/4), edit this section to select the interface to use for HA1 communication.
4. Edit the Data Link (HA2) to use **Port** ethernet 1/3 and add the IP address of this peer and the **Gateway** IP address for the subnet.

STEP 5 | **Commit** the changes.

STEP 6 | Set up the passive HA peer within the same Azure Resource Group.

1. Deploy the second instance of the firewall.
   - Download the custom template and parameters file from GitHub.
   - Log in to the Azure Portal.
   - Search for **custom template** and select **Deploy from a custom template**.
   - Select **Build your own template in the editor** > **Load file**.
   - Select the **azuredeploy.json** that you downloaded earlier, and **Save**.
   - Complete the inputs, agree to the terms and **Purchase**.

     Make sure to match the following inputs to that of the firewall instance you have already deployed— Azure subscription, name of the Resource Group, location of the Resource Group, name of the existing VNet into which you want to deploy the

 ©2024 Palo Alto Networks, Inc.

firewall, VNet CIDR, Subnet names, Subnet CIDRs, and start the IP address for the management, trust and untrust subnets.

2. Repeat Step 1and Step 2to set up the interfaces and configure the firewall as the passive HA peer.
3. Skip Step 3 and complete Enable HA (Step 5). In Step 4 modify the IP addresses as appropriate for this passive HA peer.

STEP 7 | After you finish configuring both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls, and view the High Availability widget.
2. On the active firewall, click the **Sync to peer** link.
3. Confirm that the firewalls are paired and synced, as shown as follows:
   - On the passive firewall: the state of the local firewall should display **passive** and the **Running Config** should show as **synchronized**.
   - On the active firewall: The state of the local firewall should display **active** and the **Running Config** should show as **synchronized**.
4. On the passive peer, verify that the VM-Series plugin configuration is now synced.

   Select **Device** > **VM-Series** and validate that you can view the Azure HA configuration that you had omitted configuring on the passive peer.

 ©2024 Palo Alto Networks, Inc.

Use the ARM Template to Deploy the VM-Series Firewall

In addition to Marketplace based deployments, Palo Alto Networks provides a GitHub repository which hosts sample ARM templates that you can download and customize for your needs. ARM templates are JSON files that describe the resources required for individual resources such as network interfaces, a complete virtual machine or even an entire application stack with multiple virtual machines.

ARM templates are for advanced users, and Palo Alto Networks provides the ARM template under the community supported policy. To learn about ARM templates, refer to the Microsoft documentation on ARM Templates.

To simplify the deployment of all the required resources, the two-tier sample template (https://github.com/PaloAltoNetworks/azure/tree/master/two-tier-sample) includes two json files:

- **Template File**—The azureDeploy.json is the main resources file that deploys all the components within the resource group.
- **Parameters File**—The azureDeploy.parameters.json is the file that includes the parameters required to successfully deploy the VM-Series firewall in the VNet. It includes details such as the virtual machine tier and size, username and password for the firewall, the name of the storage container for the firewall. You can customize this file for your Azure VNet deployment.

To help you deploy the firewall as a gateway for Internet-facing applications, the template provisions the VM-Series firewall, a database server, and a web server. The VNet uses the private non-routable IP address space 192.168.0.0/16. You can modify the template to use 172.16.0.0/12, or 10.0.0.0/8.

The ARM template also provides the necessary user-defined rules and IP forwarding flags to enable the VM-Series firewall to secure the Azure resource group. For the five subnets—Trust, Untrust, Web, DB, and NAT—included in the template, you have five route tables, one for each subnet with user defined rules for routing traffic to the VM-Series firewall and the NAT virtual machine.

For the four subnets—Trust, Untrust, Web, and DB—included in the template, you have four route tables, one for each subnet with user defined rules for routing traffic to the VM-Series firewall.

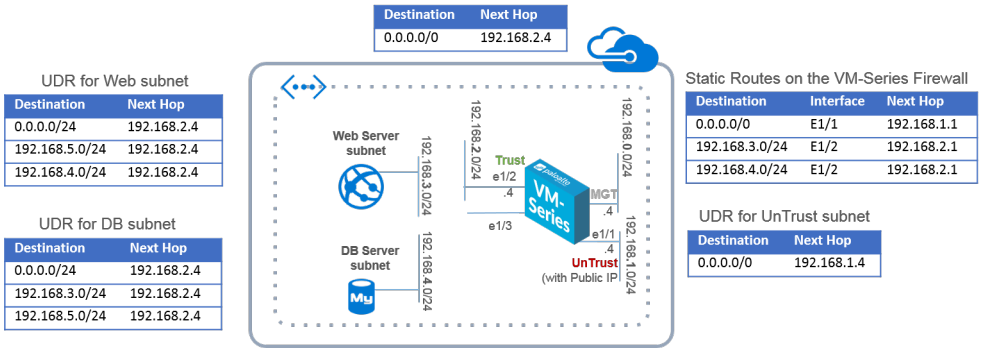

Figure 4: Deploying VM-Series Firewall using the ARM Template

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Download the two-tier sample ARM template from the GitHub repository.

Download and save the files to a local client: https://github.com/PaloAltoNetworks/azure/tree/master/two-tier-sample

STEP 2 | Create a Resource Group on Azure.

1. Log in to the Azure CLI using the command: **azurelogin**

   If you need help, refer to the Azure documentation on installing the CLI, or for details on how to access the CLI on Azure Government or Azure China.
2. Switch to Resource Manager mode using the command:**azureconfig mode arm**
3. Create a resource group.

STEP 3 | Deploy the ARM template.

1. Open the Parameters File with a text editor and modify the values for your deployment:

   In Azure China, you must edit the path for the storage account that hosts the VHD image required to deploy the VM-Series firewall. In the variables section of the template file, find the parameter called ***userImageNameURI*** *and replace the value with the location where you saved the VHD image.*
2. Deploy the template in the resource group you created.

```
azure group create -v -n “<YourResourceGroupName>”

-l “<YourAzureLocation>” -d “<GiveASmallDeploymentLabel>”

-f azureDeploy.json -e azureDeploy.parameters.json
```

3. Check the progress/status of the deployment from the Azure CLI:

```
azure group deployment show "<YourResourceGroupName>"
```

 ©2024 Palo Alto Networks, Inc.

```
“<YourDeploymentLabel>“
```

When the template is successfully deployed the `ProvisioningStateis Running`.

*If the `ProvisioningStateis Failed`, you must check for errors on the Azure portal at **Resource Group** > **Events**. Filter for only events in the last one hour, select the most recent events, and drill down to find the errors.*

4. Verify that you have successfully deployed the VM-Series firewall.
   1. Select **Dashboard** > **Resource Groups**, select the resource group.
   2. Select **All Settings** > **Deployments** > **Deployment History** for detailed status.

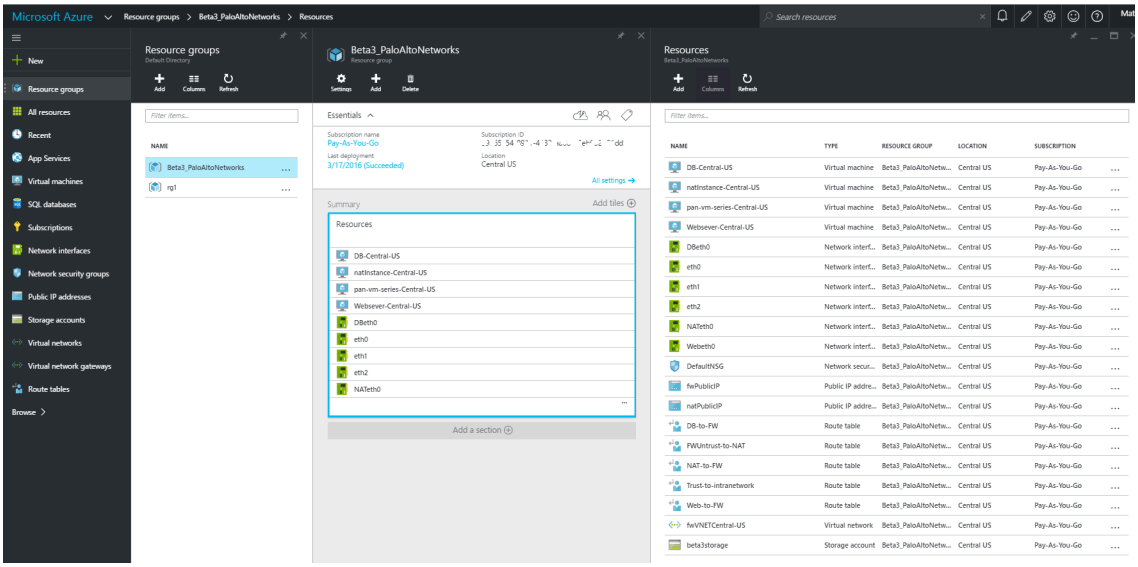

The address space within the VNet uses the prefix 192.168, which is defined in the ARM template.

5. Attach a public IP address to the untrust interface on the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Configure the firewall as a VNet gateway to protect your Internet-facing deployment.

1. Log in to the management interface IP address on the firewall.
2. Configure the dataplane network interfaces as Layer 3 interfaces on the firewall (**Network** > **Interfaces** > **Ethernet**).
3. Add static rules to the virtual router on the firewall. To route traffic through the firewall in this example, you need three static routes on the firewall (**Network** > **Virtual Routers**, select the router and click **Static Routes**):
    1. Route all outbound traffic through the UnTrust zone, ethernet1/1 to the Azure router at 192.168.1.1.
    2. Route all inbound traffic destined to the web server subnet through the Trust zone, ethernet1/2 to the Azure router at 192.168.2.1.
    3. Route all inbound traffic destined to the database server subnet through the Trust zone, ethernet1/2 to the Azure router at 192.168.2.1.
4. Create security policy rules (**Policies** > **Security**) to allow inbound and outbound traffic on the firewall. You also need security policy rules to allow appropriate traffic from the web server subnet to the database server subnet and vice versa.
5. **Commit** the changes on the firewall.
6. Verify that the VM-Series firewall is securing traffic (**Monitor** > **Logs** > **Traffic**).

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series and Azure Application Gateway Template

The VM-Series and Azure Application Gateway template is a starter kit that you can use to deploy VM-Series firewalls to secure web workloads for internet-facing deployments on Microsoft Azure (currently not available for Azure China).

This template deploys two VM-Series firewalls between a pair of (external and internal) Azure load balancers. The external load balancer is an Azure Application Gateway, which is an HTTP (Layer 7) load balancer that also serves as the internet-facing gateway, which receives traffic and distributes it through the VM-Series firewall on to the internal load balancer. The internal load balancer is an Azure Load Balancer (Layer 4) that fronts a pair of web servers. The template supports the BYOL and the Azure Marketplace versions of the VM-Series firewall.

As demand on your web workloads increases and you increase capacity for the web server tier you can manually deploy additional VM-Series firewalls to secure your web server tier.

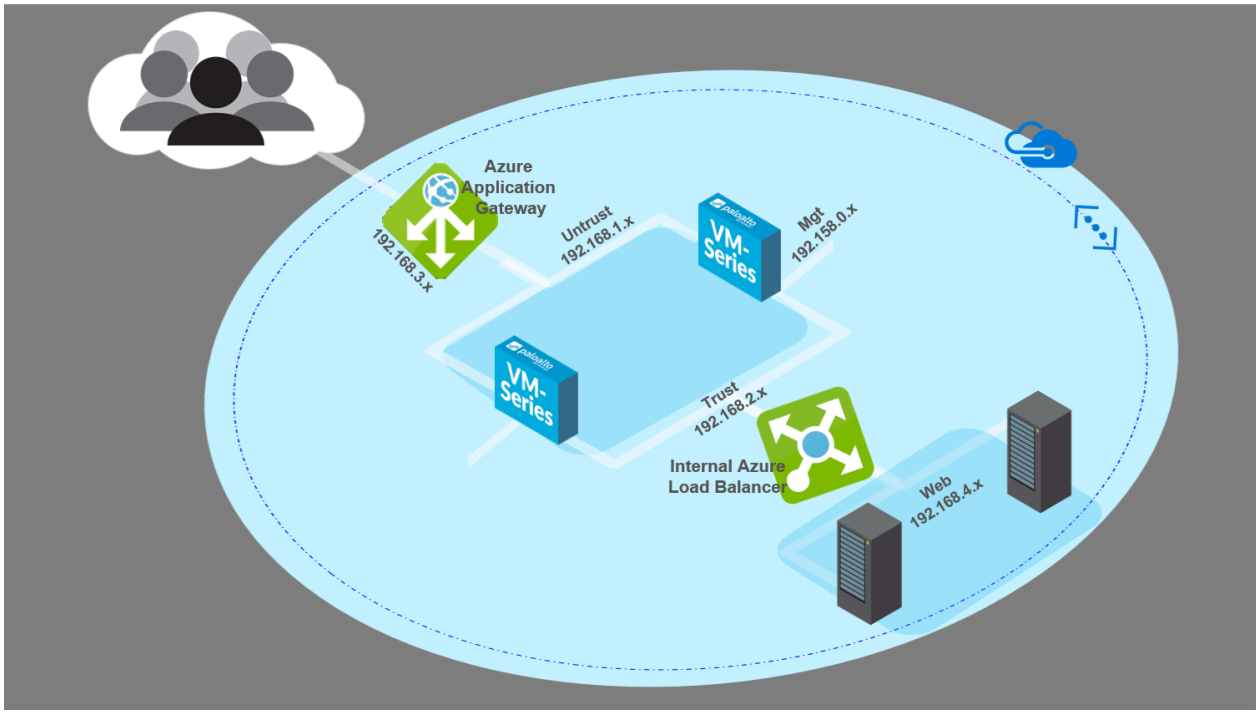

- VM-Series and Azure Application Gateway Template
- Start Using the VM-Series & Azure Application Gateway Template

 ©2024 Palo Alto Networks, Inc.

VM-Series and Azure Application Gateway Template

The VM-Series and Azure Application Gateway template launches an Azure Application Gateway (Layer 7 load balancer) and an Azure (Layer 4) load balancer. Nested between the Application gateway and the load balancer are a pair of VM-Series firewalls in an Availability Set, and a pair of sample web servers running Apache2 on Ubuntu in another Availability Set. The Availability Sets provide protection from planned and unplanned outages. The following topology diagram shows the resources that the template deploys:

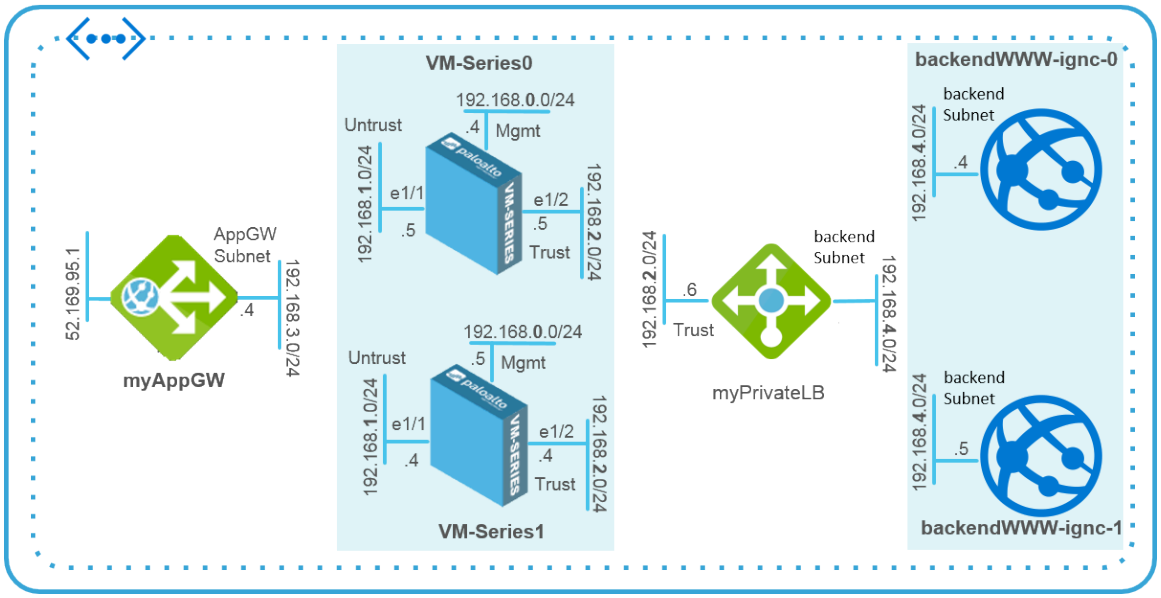

You can use a new or an existing storage account and resource group in which to deploy all the resources for this solution within an Azure location. It does not provide default values for the resource group name and storage account name, you must enter a name for them. While you can create a new or use an existing VNet, the template creates a default VNet named *vnet-FW* with the CIDR block 192.168.0.0/16, and allocates five subnets (192.168.1.0/24 - 192.168.5.0/24) for deploying the Azure Application Gateway, the VM-Series firewalls, the Azure load balancer and the web servers. Each VM-Series firewall is deployed with three network interfaces—ethernet0/1 in Mgmt subnet (192.168.0.0/24), ethernet1/1 in Untrust subnet (192.168.1.0/24), and ethernet1/2 in the Trust subnet (192.168.2.0/24).

The template creates a Network Security Group (NSG) that allows inbound traffic from any source IP address on ports 80,443, and 22. It also deploys the pair of VM-Series firewalls and the web server pair in their respective Availability Sets to ensure that at least one instance of each is available during a planned or unplanned maintenance window. Each Availability Set is configured to use three fault domains and five update domains.

The Azure Application Gateway acts as a reverse-proxy service, which terminates a client connection and forwards the requests to back-end web servers. The Azure Application Gateway is set up with an HTTP listener and uses a default health probe to test that the VM-Series firewall IP address (for ethernet1/1) is healthy and can receive traffic.

The template does not provide an auto-scaling solution; you must plan your capacity needs and then deploy additional resources to Adapt the Template *for your deployment.*

 ©2024 Palo Alto Networks, Inc.

The VM-Series firewalls are not configured to receive and secure web traffic destined to the web servers. Therefore, at a minimum, you must configure the firewall with a static route to send traffic from the VM-Series firewalls to the default router, configure destination NAT policy to send traffic back to the IP address of the load balancer, and configure Security policy rules. The NAT policy rule is also required for the firewall to send responses back to the health probes from the HTTP listener on the Azure Application Gateway. To assist you with a basic firewall configuration, the GitHub repository includes a sample configuration file called *appgw-sample.xml* that you can use to get started.

Start Using the VM-Series & Azure Application Gateway Template

The VM-Series & Azure Application Gateway template launches all the resources you need to deploy and secure your web workloads for Internet facing deployments on Microsoft Azure, excluding Azure China. This section provides details on how to deploy the template, configure the firewalls to route and secure traffic destined to the web servers, and extend the capabilities and resources that this template provides to accommodate your deployment needs.

- Deploy the Template to Azure
- VM-Series and Azure Application Gateway Template Parameters
- Sample Configuration File
- Adapt the Template

Deploy the Template to Azure

Use the following instructions to deploy the template to Azure.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Deploy the template.

Currently not available for deploying in Azure China.

1. Access the template from https://github.com/PaloAltoNetworks/azure-applicationgateway
2. Click **Deploy to Azure**.
3. Fill in the details for deploying the template. See VM-Series and Azure Application Gateway Template Parameters for a description and the default values, if any, for each parameter.

   At a minimum, you have to pick the **Azure Subscription**, **Resource Group**, **Location**, **Storage Account Name**, and a **Username/password** or **SSH Key** for the administrative account on the VM-Series firewalls.
4. Click **Purchase** to accept the terms and conditions and deploy the resources.

   If you have validation errors, click to view the details and fix your errors.
5. On the Azure portal, verify that you have successfully deployed the template resources, including the VM-Series firewalls.
   1. Select **Dashboard** > **Resource Groups**, select the resource group.
   2. Select **Overview** to review all the resources that have been deployed. The deployment status should display **Succeeded**.

 ©2024 Palo Alto Networks, Inc.

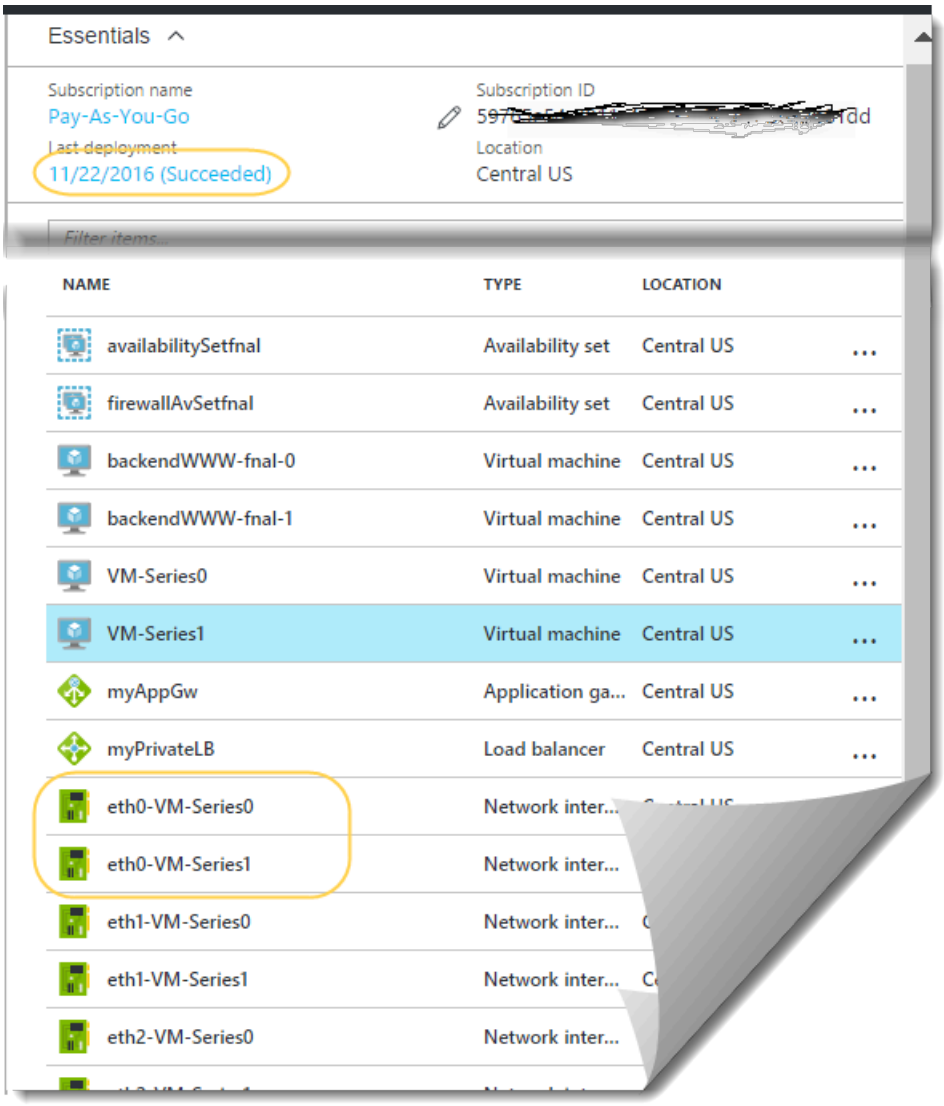

3. Note the Public IP address or the DNS name assigned to **eth0-VM-Series0** and **eth0-VM-Series1** to access the management interface of the VM-Series firewalls.

STEP 2 | Log in to the firewalls.

1. Using a secure connection (https) from your web browser, log in to the IP address for eth0-VM-Series0 or the DNS name for the firewall.
2. Enter the username/password you defined in the parameters file. You will see a certificate warning; that is okay. Continue to the web page.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure the VM-Series firewall.

You can either configure the firewall manually or import the Sample Configuration File provided in the GitHub repository and customize it for your security needs.

- **Configure the firewall manually**—You must do the following at a minimum:
  1. Configure the dataplane network interfaces as Layer 3 interfaces on the firewall (**Network** > **Interfaces** > **Ethernet**).
  2. Add a static rule to the virtual router on the firewall. This static rule specifies the firewall's untrust interface IP address as the nexthop address for any traffic destined for ethernet1/1. (**Network** > **Virtual Routers**, select the router and click **Static Routes**).
  3. Create security policy rules (**Policies** > **Security**) to allow inbound and outbound traffic on the firewall.
  4. Add NAT policies (**Policies** > **NAT**). You must create destination NAT and source NAT rules on the firewall to send traffic to the web servers and back out to the client who initiated the request.

     The destination NAT rule is for all traffic that arrives on the firewall's untrust interface. This rule is required to translate the destination IP address on the packet to that of the internal load balancer so that all traffic is directed to the internal load balancer and on to the backend web servers.

     The source NAT rule is for all traffic from the backend web server and destined to the untrust interface on the firewall. This rule translates the source address to the IP address of the trust interface on the firewall
  5. **Commit** your changes.

     - **Import the sample configuration file:**
  6. Download and save the Sample Configuration File to your local client.
  7. Select **Device** > **Setup** > **Operations**, click **Import named configuration snapshot**, **Browse** to the sample configuration file that you have saved locally, and click **OK**.
  8. Click **Load named configuration snapshot**, select the **Name** of the sample configuration file you just imported, and click **OK**.
  9. Change the IP address of the address objects and the static route to match the IP address from the CIDR block you used. Update address objects to use the private IP addresses for eth1-VM-Series0 and eth1-VM-Series1.
  10. Important! Create a new admin user account. Select **Device** > **Administrators** and **Add** a new account.
  11. Modify the **Hostname** in the General Settings widget in **Device** > **Setup** > **Management**.
  12. **Commit** your changes, and log out. The commit overwrites the running configuration with the sample configuration file and updates you just made. On commit, the hostname and the administrator user account that you specified when deploying the template are overwritten. You will now need to log in using the new admin user account and password.

      - **Log in to the firewall**—Use the credentials you created and delete the pandemo administrative account imported as part of the sample configuration file.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Log in and configure the other instance of the VM-Series firewall.

See step Configure the VM-Series firewall.

STEP 5 | Verify that you have configured the firewalls properly.

From your web browser, use http to access the IP address or DNS name for the app gateway. You should be able to view the default Apache 2 Ubuntu web page.

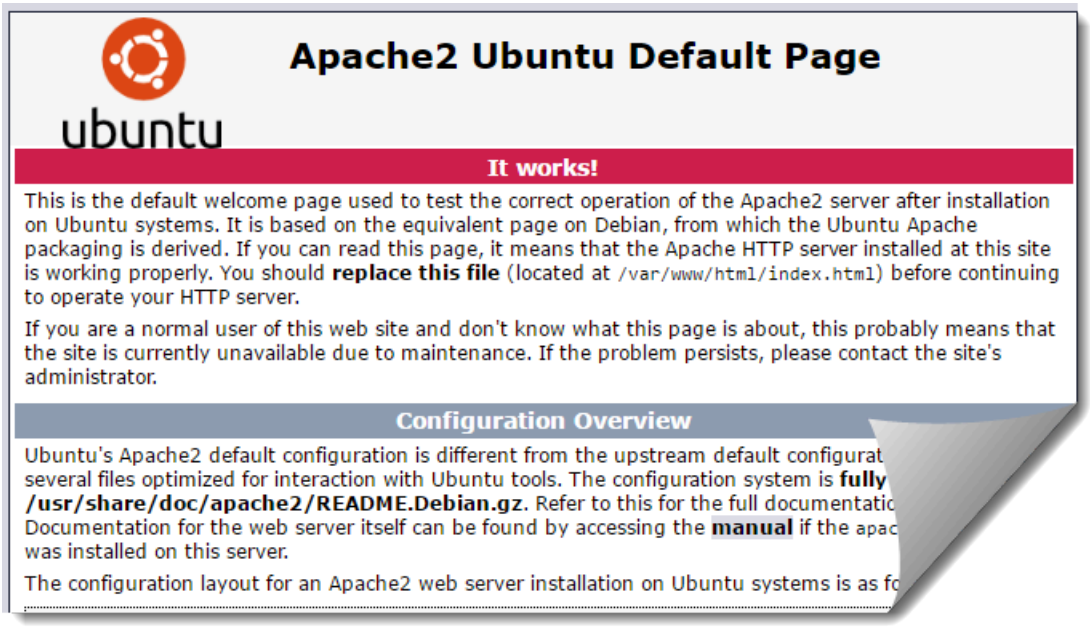

If you have used the sample configuration firewall, log in to the firewall and view the Traffic logs generated on session start in **Monitor** > **Logs** > **Traffic**.

VM-Series and Azure Application Gateway Template Parameters

The following table lists the required and optional parameters and the default values, if any.

Parameter	Description
Resource group	Create new or use existing (no default).
Subscription	The type of Azure subscription you will use to cover the cost of the resources deployed with the template.
Location	Select the Azure location to which you want to deploy the template (no default).
Network Security Group	
Network Security Group Name	The network security group limits the source IP addresses from which the VM-Series firewalls and web servers can be accessed. Default: nsg-mgmt

 ©2024 Palo Alto Networks, Inc.

Parameter	Description
Network Security Group Inbound Src IP	The source IP addresses that can log in to the management port of the VMs deployed by the template. The default value 0.0.0.0/0 means you can log into the firewall management port from any IP address.
Storage Account	
Storage Account Name	Create new or enter the name of an existing Storage Account (no default). The name must be globally unique.
Storage Account Type	Choose between standard and premium storage and your data replication needs for local redundancy, geo-redundancy, and read-access geo-redundancy. The default option is Locally Redundant Storage (LRS). The other options are Standard GRS, Premium LRS, and Standard RAGRS.
VNet	
Virtual Network	Create new or enter the name of an existing VNet. The default name for the VNet is vnet-FW
Virtual Network Address Prefix	192.168.0.0/16
Azure Application Gateway	
App Gateway Name	myAppGw
App Gateway DNS Name	Enter a globally unique DNS name for the Azure Application Gateway.
App Gateway Subnet Name and Prefix	Default name is AppGWSubnet and the subnet prefix is 192.168.3.0/24.
Azure Load Balancer and Web Servers	
Internal Load Balancer Name	myPrivateLB
Internal Load Balancer Subnet Name and Prefix	Default name is backendSubnet and the subnet prefix is 192.168.4.0/24.
Backend Vm Size	The default size is Standard tier D1 Azure VM. Use the drop-down in the template to view the other Azure VM options available for the backend web servers.

 ©2024 Palo Alto Networks, Inc.

Parameter	Description
Firewalls	
Firewall Model	Choose from BYOL or PAYG (bundle 1 or bundle 2, each bundle includes the VM-300 and a set of subscriptions).
Firewall Vm Name and Size	The default name for the firewall is VM-Series, and the default size is Standard tier D3 Azure VM. Use the drop-down in the template to view the other Azure VM options available for the VM-Series firewalls
Mgmt Subnet Name and Prefix	The management subnet for the VM-Series firewalls and the web servers deployed in this solution. Default name is Mgmt and the subnet prefix is 192.168.0.0/24.
Mgmt Public IP Address Name	Enter a hostname to access the management interface on each firewall. The names must be globally unique.
Trusted Subnet Name and Prefix	The subnet to which eth1/1 on the VM-Series firewall is connected; this subnet connects the VM-Series firewall to the Azure Application gateway. The firewall receives web traffic destined to the web servers on eth1/1. Default name is Trust and the subnet prefix is 192.168.2.0/24.
Untrusted Subnet Name	The subnet to which eth1/2 on the VM-Series firewall is connected. The firewall receives return and outbound web traffic on this interface. Default name is Untrust and the subnet prefix is 192.168.1.0/24. The name must be globally unique.
Username	Enter the username for the administrative account on the VM-Series firewalls and the web servers.
Authentication Type	You must either enter a password for authentication or use an SSH public key (no default).

Sample Configuration File

To help you get started, the GitHub repository contains a sample configuration file named *appgw-sample.xml* that includes the following rules/objects:

- **Address objects**—Two address objects, `firewall-untrust-IP` and `internal-load-balancer-IP`, which you will need to modify to match the IP addresses in your setup. You need to modify these address objects to use the private IP addresses assigned to eth1-VM-Series0 and eth1-VM-Series1 on the Azure portal.

 ©2024 Palo Alto Networks, Inc.

- **Static route**—The default virtual router on the firewall has a static route to 192.168.1.1, and this IP address is accurate if you use the default template values. If you have changed the Untrust subnet CIDR, you'll need to update the IP address to match your setup. All traffic coming from the backend web servers, destined for the application gateway, uses this IP address as the next hop for delivering packets to the untrust interface on the firewall.
- **NAT Policy Rule**—The NAT policy rule enables destination NAT and source NAT.
  - The destination NAT rule is for all traffic that arrives on the firewall's untrust interface (ethernet1/2), which is the firewall-untrust-IP address object. This rule translates the destination IP address on the packet to that of the internal load balancer so that all traffic is directed to the internal load balancer and thus to the backend web servers.
  - The source NAT rule is for all traffic from the backend web server and destined to the untrust network interface on the firewall. This rule translates the source address to the IP address of the trust interface on the firewall (ethernet1/2).
- **Security Policy Rule**—Two Security policy rules are defined in the sample configuration file. The first rule allows all inbound web-browsing traffic and generates a log at the start of a session on the firewall. The second rule blocks all other traffic and generates a log at the start and end of a session on the firewall. You can use these logs to monitor all traffic to the web servers in this deployment.
- **Administrative User Credentials**— The sample configuration file includes a username and password for logging in to the firewall, which is set to pandemo/demopassword. After you import the sample configuration, you must either change the password and set it to a strong, custom password or create a new administrator account and delete the pandemo account.

Adapt the Template

As your needs evolve, you can scope your capacity needs and extend the template for your deployment scenario. Here are some ways you can build on the starter template to meet your planned capacity needs:

- Deploy additional VM-Series firewalls behind the Azure Application Gateway. You can manually install more VM-Series firewalls into the same Availability Set or launch a new Availability Set and manually deploy additional VM-Seri es firewalls.
- Configure the VM-Series firewalls beyond the basic configuration provided in the sample configuration file in the GitHub repository.
- Enable HTTPS load balancing (SSL offload) on the Azure Application Gateway. Refer to the Azure documentation for details.
- Add or replace the sample web servers included with the template.

 ©2024 Palo Alto Networks, Inc.

Secure Kubernetes Services on Azure

To secure Azure Kubernetes services, you must first install the Azure plugin on Panorama and configure an Azure Secure Kubernetes Services on Azure deployment. The Azure plugin for Panorama supports tag-based VM monitoring and Secure Kubernetes Services on Azure, secures inbound traffic for Azure Kubernetes Services (AKS) clusters, and monitors outbound traffic from AKS clusters. The Panorama orchestrated deployment allows you to leverage Azure auto scale metrics and the scale-in and scale-out thresholds to manage surges in demand for application workload resources by independently scaling the VM-Series firewalls.

To secure inbound traffic for your AKS cluster, you must first Secure Kubernetes Services on Azure. The Panorama orchestrated deployment works with Secure Kubernetes Services on Azure to gather information about your network and resources, then create an auto-scaling tier of VM-Series firewalls for either Secure Kubernetes Services on Azure deployments. See the Palo Alto Networks Compatibility Matrix, to verify the minimum OS, plugin, and template versions required to secure AKS clusters.

Palo Alto Networks provides an AKS template that deploys an Azure Kubernetes Service (AKS) cluster in a new Azure VNet. The Azure plugin on Panorama helps you set up a connection which can monitor Azure Kubernetes cluster workloads, harvesting services you have annotated as “internal load balancer” and creating tags you can use in dynamic address groups. You can leverage Panorama dynamic address groups to apply security policy on inbound traffic routed to services running on your AKS cluster.

- How Does the Panorama Plugin for Azure Secure Kubernetes Services?
- Secure an AKS Cluster

How Does the Panorama Plugin for Azure Secure Kubernetes Services?

You can use VM-Series firewalls to secure inbound traffic for Azure Kubernetes Service (AKS) clusters. The VM-Series firewall can only secure services exposed by a load balancer (such as an Azure Load Balancer). Outbound traffic can only be monitored.

This chapter reviews different components that enable the Azure Plugin for Panorama to connect to an AKS cluster.

- Requirements
- A Sample Hub-and-Spoke Topology to Secure AKS Clusters
- User-Defined Routing
- AKS Cluster Communication
- Dynamic Address Groups with Kubernetes Labels

Requirements

This solution requires the following components. See the Panorama plugin information in the Compatibility Matrix for the minimum version requirements.

- VM-Series firewalls.

 ©2024 Palo Alto Networks, Inc.

- Panorama—Your Panorama version must be the same or higher than your VM-Series PAN-OS version.
- Panorama Plugin for Azure.
- A Panorama orchestrated deployment.
- Azure AKS template version 1.0. This template creates an AKS cluster.

  You must enable AKS advanced networking (CNI) for the cluster.

  An AKS deployment requires advanced networking to configure VNet Peering for the hub and spoke VNets (see A Sample Hub-and-Spoke Topology to Secure AKS Clusters).

A Sample Hub-and-Spoke Topology to Secure AKS Clusters

The following diagram illustrates a sample auto scale deployment that secures inbound traffic for Azure AKS clusters. Let's review some of the components.

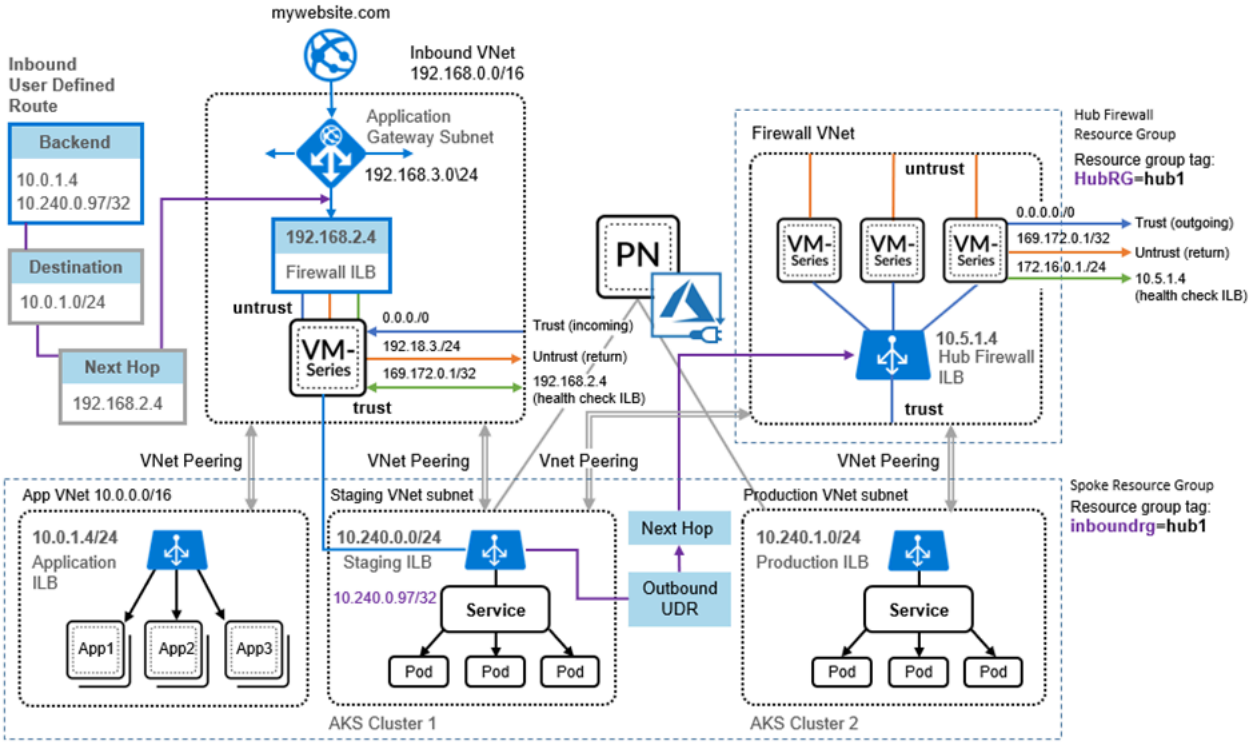

- **Auto Scaling Infrastructure**—The Azure Auto Scaling templates create the messaging infrastructure and the basic hub and spoke architecture.
- **AKS Clusters**—The Palo Alto Networks AKS template creates an AKS cluster in a new VNet. Given the name of the Spoke resource group, the template tags the VNet and AKS cluster with the Spoke resource group name, so the resource group can be discovered by the Azure Auto Scaling plugin for Panorama. The Azure plugin for Panorama queries service IP addresses on the Staging ILB to learn about AKS cluster services.

  Only one Spoke firewall scale set can be associated with an AKS Cluster; if you expose multiple services in a single AKS cluster, they must be protected by the same Spoke.

  For each resource group, create a subnet-based address group. In the above diagram, for example, create an address group for 10.240.0.0/24 (AKS Cluster 1).

 ©2024 Palo Alto Networks, Inc.

- **VNet Peering**—You must manually configure VNET peering to communicate with other VNets in the same region.

  Cross-region peering is not supported.

  You can use other automation tools to deploy AKS clusters. If you deploy in an existing VNet (the Hub Firewall VNet, for example) you must manually configure VNet peering to the inbound and outbound hub and spoke resource groups, and manually tag the VNet and AKS cluster with the resource group name.
- **User Defined Routes and Rules**—You must manually configure user-defined routes and rules (see User-Defined Routing). In the diagram above, incoming traffic can be redirected, according to UDR rules, to the Firewall ILB for inspection. Outbound traffic exiting an AKS Cluster is redirected to the Hub Firewall ILB with Azure user-defined routing (UDR) rules. The solution assumes Allow All as a default policy for Kubernetes orchestration to function as-is, but to apply policy you can use an allowlist or a denylist to allow or deny outbound traffic.

User-Defined Routing

You must manually create user-defined routing and routing rules to govern inbound or outbound traffic.

Inbound

In the above diagram, inbound traffic from the Application gateway is driven to the backend pool, and based on UDR rules, redirected to the Firewall ILB. For example, create a UDR pointing to the VNet subnet so that the traffic for Kubernetes services is pointed to the firewall ILB.

Outbound

On the Hub firewall set, for each AKS cluster being protected, you must create static routes for the cluster subnet CIDR, with the next hop being the gateway address of the Hub VNet trust subnet.

All outbound traffic for an AKS cluster is directed to the Hub firewall set with a single UDR rule.

AKS Cluster Communication

The Panorama plugin for Azure can only communicate with the AKS controller node for a given AKS cluster. For Outbound AKS traffic, the next hop is the Hub Firewall ILB. Because Outbound traffic is monitored, you must Allow All traffic. The following topics emphasize common practices that help you establish connectivity. Keep them in mind when you plan your networks and subnets.

- Create AKS Cluster Authentication
- Use An Address Group to Identify Traffic
- Add the Subnet Address Group to the Top-Level Policy
- Prevent Application Disruption when Workload and AKS Cluster VNets Are Peered

Create AKS Cluster Authentication

When you connect the AKS cluster in Azure plugin for Panorama you must enter a secret authorization token. Use Kubernetes commands to perform the following steps.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create a ClusterRole.

STEP 2 | Create a ClusterRoleBinding.

1. Create a `.yaml` file for the ClusterRoleBinding. For example, create a text file named `crb.yaml`.

```
apiVersion: rbac.authorization.k8s.io
kind: ClusterRoleBinding
metadata:
  name: default-view
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: view
subjects:
- kind: ServiceAccount
name: default
namespace: default
```

2. Use Azure Cloud Shell to apply the `crb.yaml` role binding.

   `kubectl apply -f crb.yaml`

3. View the service account you just created.

   `kubectl get serviceaccounts`

STEP 3 | Save the service account credential to a `.json` file.

1. On your local machine, change to the directory in which you want to save the credential.
2. Use kubectl commands to create the token.

   `MY_SA_TOKEN=‘kubectl get serviceaccounts default -o jsonpath=’{.secrets[0].name}’‘`

3. View the token name.

   `$ echo $MY_SA_TOKEN`

4. Display the credential.

   `kubectl get secret $MY_SA_TOKEN -o json`

You need the token when you connect the AKS cluster in Azure plugin for Panorama, in Step 3.d.

Use An Address Group to Identify Traffic

To create some granularity for monitored Outbound traffic, create an address group specifically for the AKS cluster VNet subnet (for example, 10.240.0.97/32 in the above diagram). You can then write rules that allow incoming or returning traffic rather than using Allow All.

If you create an address group, be careful to maintain the communication between the AKS controller and any worker nodes. See Add the Subnet Address Group to the Top-Level Policy.

If communication is interrupted, application traffic can be lost or your application deployment might have problems.

 ©2024 Palo Alto Networks, Inc.

Add the Subnet Address Group to the Top-Level Policy

To maintain connectivity, the address group must be part of the top-level policy in Panorama. You can configure the cluster address group, or bootstrap the cluster to configure the cluster address group.

*Add the address group to the top-level policy **before** you configure VNet peering or* User-Defined Routing.

Prevent Application Disruption when Workload and AKS Cluster VNets Are Peered

If an AKS cluster co-exists with VM workloads that run in separate VNets, and the VNet is peered with both the workload spoke (Inbound) and the Hub (Outbound), you must create address groups to differentiate the workloads and the AKS traffic, and add the address group to Top-Level Policy as described above.

Dynamic Address Groups with Kubernetes Labels

When monitoring an AKS cluster resource, the Azure plugin automatically generates the following IP tags for AKS services.

```
aks.<aks cluster name>.<aks service name>
```

Tags are not generated for nodes, pods, or other resources.

If the AKS service has any labels, the tag is as follows (one per label):

```
aks.<aks cluster name>.svc.<label>.<value>
```

If a labelSelector tag is defined for a cluster, the plugin generates the following IP tag:

```
aks_<labelSelector>.<aks cluster name>.<aks service name>
```

Secure an AKS Cluster

To enable Panorama to connect to the load balancers in an Azure Kubernetes Services (AKS) cluster, you must enable the Azure plugin on Panorama to establish a connection with your AKS cluster. Then, you must configure the device groups and templates to which the firewalls belong so that Panorama can push configuration objects and policy rules to your managed firewalls.

- Before You Begin
- Use the Template to Deploy an AKS Cluster
- Connect the AKS Cluster in Azure Plugin for Panorama
- Set Up VNet Peering
- Redirect Traffic to a Firewall ILB
- Apply Policy to Relevant AKS Service
- Deploy and Secure AKS Services

 ©2024 Palo Alto Networks, Inc.

Before You Begin

To secure AKS you must first deploy the Azure Auto Scaling solution available on GitHub.

To secure a web application running as a service within a Kubernetes cluster you must plan the VNets, subnets, and UDRs. VM-Series firewalls and Panorama provide you security and visibility of your Kubernetes services.

- Review "How Does the Panorama Plugin for Azure Secure Kubernetes Services?".
- You must have AKS advanced networking to use the Palo Alto Networks AKS template.
- Design your AKS subnets **before** you deploy AKS clusters. Review A Sample Hub-and-Spoke Topology to Secure AKS Clusters, and AKS Cluster Communication.
  - The template creates a single AKS cluster (Service) as a sample. You must specify CIDR ranges for the VNet, VNet subnet, and the service. The CIDR ranges must not overlap
  - Size your subnets to your requirements. Avoid unnecessarily large ranges, as they can affect performance.
  - See User-Defined Routing. Specify specific UDR routes rather than broad subnet-specific routes.
- Plan how you want to peer your VNets. If you are peering AKS clusters, be sure you have read AKS Cluster Communication.
- Think about the ways in which you want to identify traffic.
  - If you plan to use an address group on Outbound AKS traffic, see Add the Subnet Address Group to the Top-Level Policy.
  - If you have service names or tags that are not unique across namespaces, use the label selector to filter both a tag and a namespace so that you get a unique result.

Use the Template to Deploy an AKS Cluster

The Azure AKS template is a sample that provisions a cluster in a new VNet.

STEP 1 | On GitHub, go to PaloAltoNetworks/azure-aks and locate the build package in the repository.

STEP 2 | Unzip the build package. Edit the files `azuredeploy.json` and `parameters.json` for your own deployment, and save.

STEP 3 | Issue the following Azure CLI commands to deploy the template.

```
az group deployment validate --resource-group RG_NAME
--template-file azuredeploy.json
--parameters @parameters.json
```

```
az group deployment create --name DEPLOYMENT_NAME
--resource-group RG_NAME
--template-file azuredeploy.json
--parameters @parameters.json
```

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Deploy your applications or services on the AKS Cluster.

1. Annotate your service YAML file so that the type is load balancer, and annotate it as service.beta.kubernetes.io/azure-load-balancer-internal: "true". For example:

```
 apiVersion: v1
   kind: Service
 metadata:
  name: azure-vote-front
  labels:
    service: "azure-vote-front"
    tier: "stagingapp"
  annotations:
    service.beta.kubernetes.io/azure-load-balancer-internal:
"true"
 spec:
  type: LoadBalancer
ports:
- port: 80
selector:
 app: azure-vote-front
```

2. If you have not done so, create AKS cluster authentication before continuing.
3. Deploy your service on your AKS cluster.

   For example, you can deploy your application through kubectl:

   kubectl apply -f myapplication.yaml

   For a sample, see: https://github.com/Azure-Samples/azure-voting-app-redis/blob/master/azure-vote-all-in-one-redis.yaml
4. Use kubectl to get the service IP for the deployed service.

```
kubectl get services -o wide
```

```
NAME              TYPE           CLUSTER-IP    EXTERNAL-IP   PORT(S)          AGE    SELECTOR
azure-vote-back   ClusterIP      10.0.77.21    <none>        6379/TCP         2d23h  app=azure-vote-back
azure-vote-front  LoadBalancer   10.0.18.189   10.240.0.97   80:31937/TCP     2d23h  app=azure-vote-front
kubernetes        ClusterIP      10.0.0.1      <none>        443/TCP          2d23h  <none>
```

   In the EXTERNAL-IP column 10.240.0.97 is for the ILB, according to your annotation in Step a. Use the service IP to create a user defined route on Azure.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Create a UDR rule to point your service to the Firewall ILB behind the Application Gateway.

In Azure, go to your inbound spoke resource group, view the route table and add a new route based on the destination service IP. In the following screen, the value in the tov1service **ADDRESS PREFIX** column is the service IP.

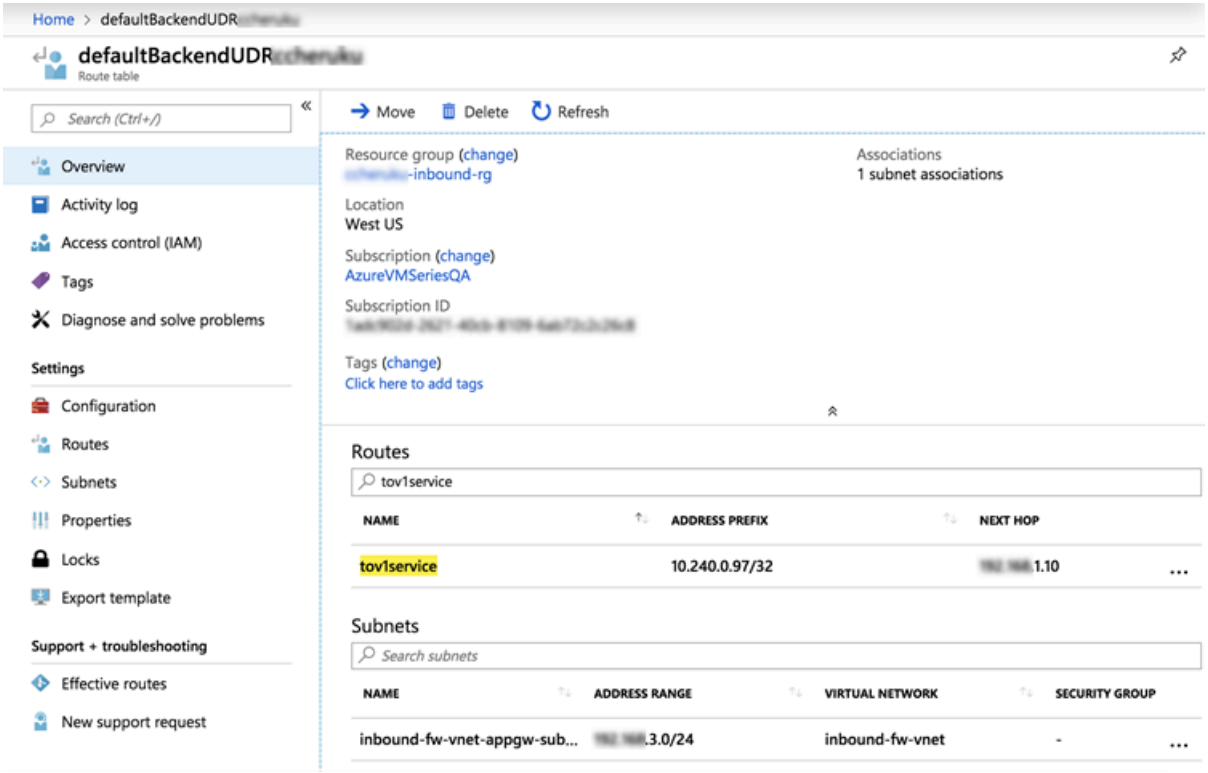

Connect the AKS Cluster in Azure Plugin for Panorama

This task assumes you have deployed a Panorama orchestrated deployment, and that you have created templates, template stacks and device groups.

See the Panorama online help for more on filling out each form.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Select **Panorama** > **Azure** > **Deployments** to view the monitoring definition you created when you configured the deployment. As shown below, if **Auto Program Routes** is enabled, the firewall routes are programmed for you.

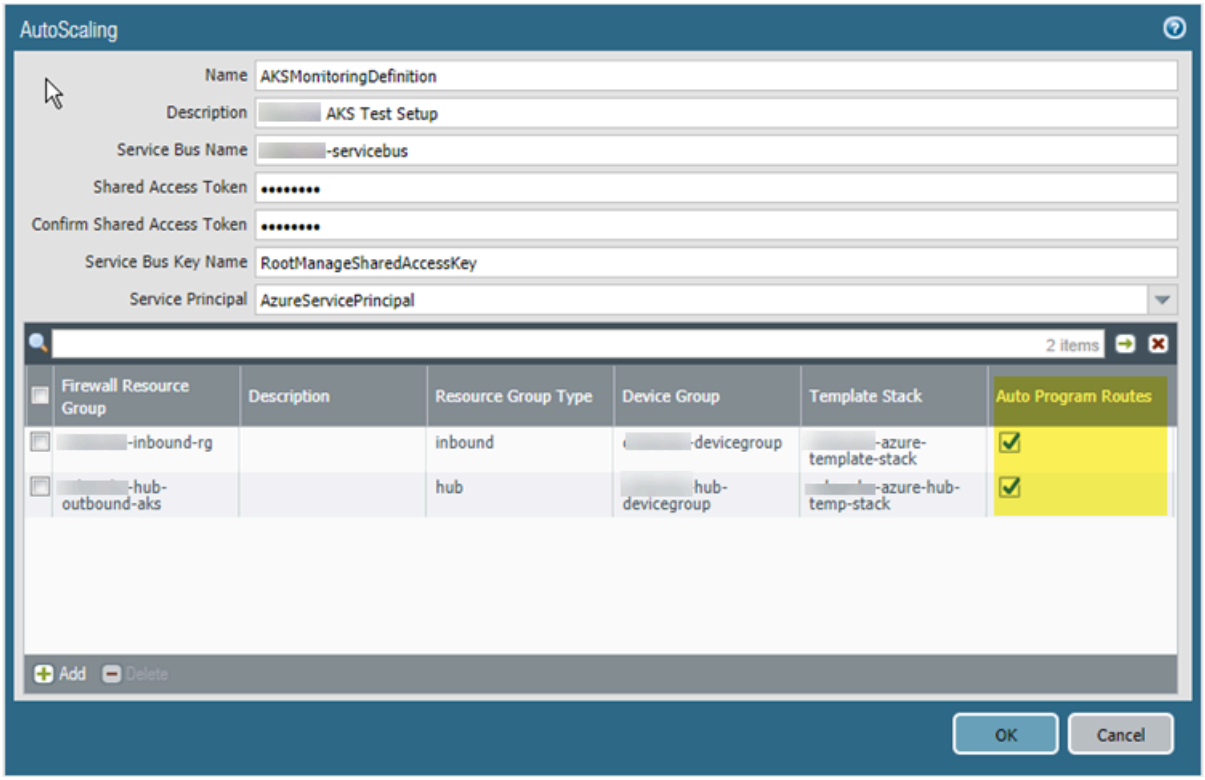

 ©2024 Palo Alto Networks, Inc.

STEP 2 | In AKS, tag your Resource Groups. The tags are name/value pairs.

1. Select **Home** > **Resource groups** and choose a resource group.
2. Select **Tags** and define name/value pairs. As shown in the following figure, the tag names must be inboundgrouprg and HubRG:
   - inboundgrouprg—your spoke resource group name
   - HubRG—your hub resource group name

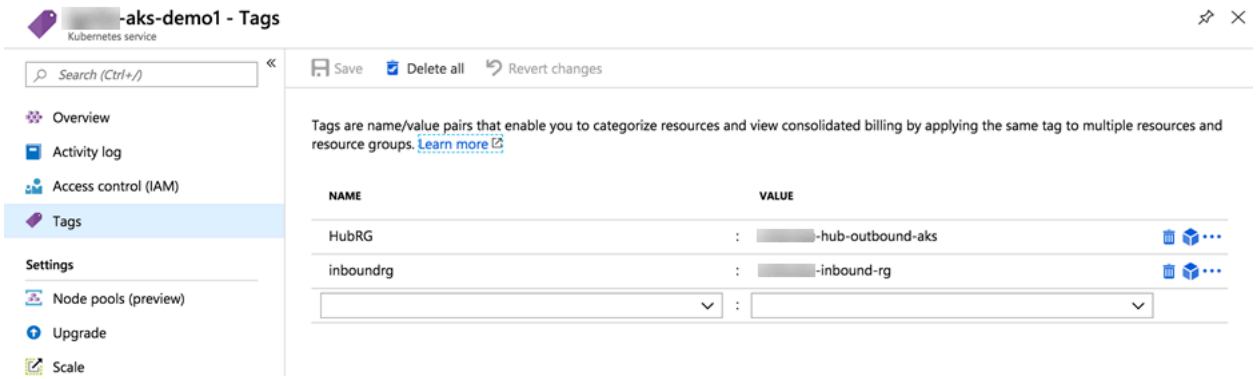

The template takes the name of the Spoke resource group as a parameter, and tags the VNet and AKS cluster with the Spoke resource group name so that it can be discovered by the Panorama plugin for Azure.

The templates deploy resources in separate VNets. If you manually deploy the AKS cluster and service in the same VNet as the Spoke firewall set, you must manually create tags for the spoke resource group name.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | In Panorama, select **Panorama** > **Azure** > **Setup.**

1. On the **General** tab, enable monitoring.
2. On the **Notify Groups** tab, **Add** a notification group and select the device groups to be notified.

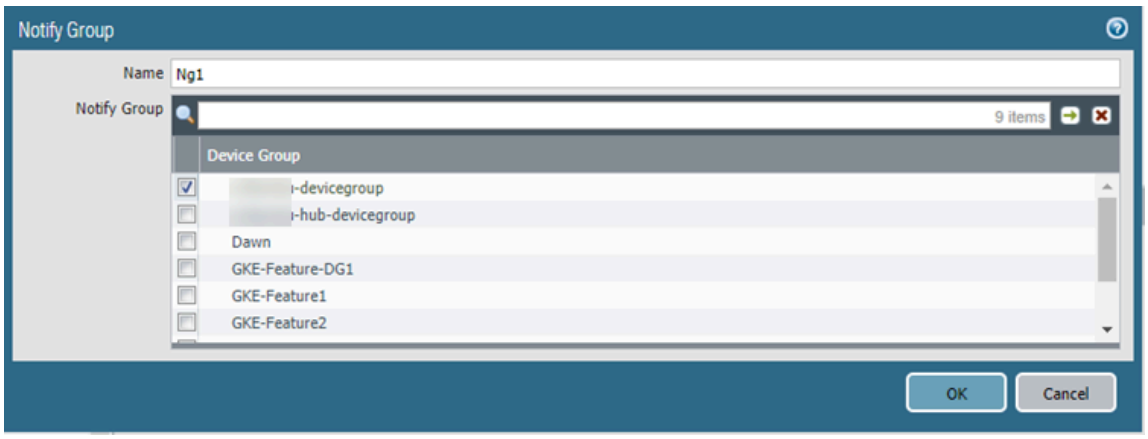

3. On the **Service Principal** tab, **Add** and **Validate** a service principal.

   Use the Service Principal you created for the orchestrated deployment.

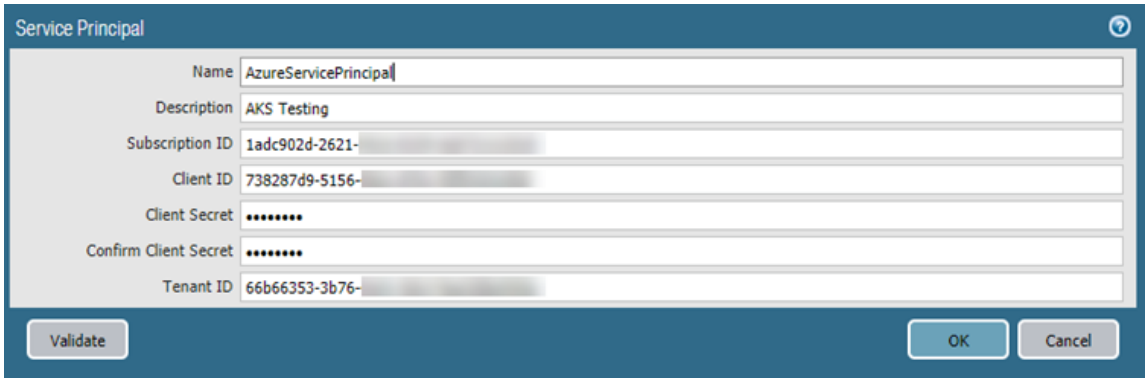

4. On the **AKS Cluster** tab, **Add** an AKS cluster.
   - Enter the exact name of the AKS cluster.
   - Enter the API server address. To find the address in Azure, view your AKS service and select Overview.
   - Upload the AKS credential JSON file (see Create AKS Cluster Authentication).
5. Fill in the remaining fields and **Add** one or more tags.

   If you have service names or tags that are not unique across namespaces, use the label selector to filter both a tag and a namespace so that you get a unique result.

 ©2024 Palo Alto Networks, Inc.

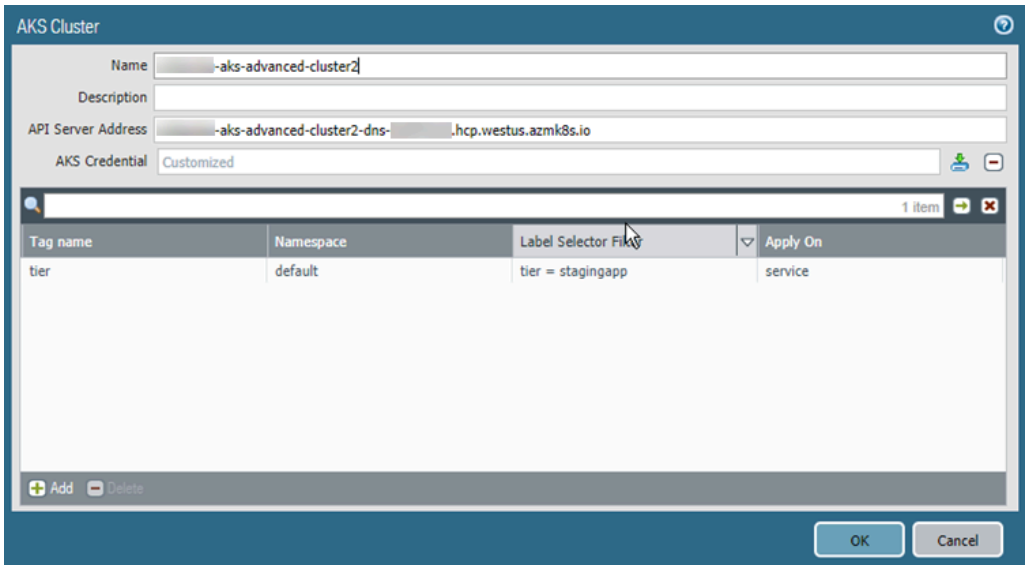

STEP 4 | Select **Panorama** > **Azure** > **Monitoring Definition**

1. Add a Monitoring definition.
2. Enter a name and description, and select **AKS Cluster Monitoring**.
3. Select an **AKS Cluster** and a **Notify Group**, check **Enable**, and click **OK**.

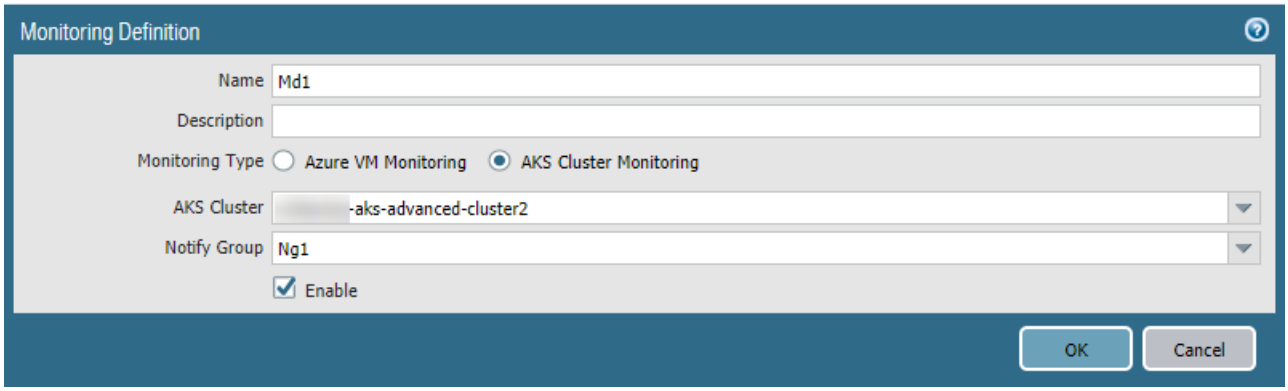

Set Up VNet Peering

If you plan to use an address group to identify traffic, be sure to add the subnet address group to your top-level Panorama plicy before you configure peering.

After deploying an AKS cluster, set up VNet Peering from the Inbound VNet to your cluster, and from your cluster to the Firewall VNet.

Redirect Traffic to a Firewall ILB

You must manually create user defined routes (UDRs) and routing rules to redirect traffic to a particular ILB. For an example, see how the diagram in "How Does the Panorama Plugin for Azure Secure Kubernetes Services?" depicts an inbound UDR.

STEP 1 | Create URL routing rules that redirect web traffic to the appropriate backend pool.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Update the UDR rules for the application gateway subnet to add a route for the service CIDR, with the next hop being the Inbound Firewall Load Balancer from the Spoke firewall resource group.

Apply Policy to Relevant AKS Service

STEP 1 | In Panorama, select Policies.

STEP 2 | In the **Device Group** list, choose the device group for your AKS service.

STEP 3 | **Add** a Security Policy rule. Fill out the form, and on the **Destination** tab **Add** the destination address or address group.

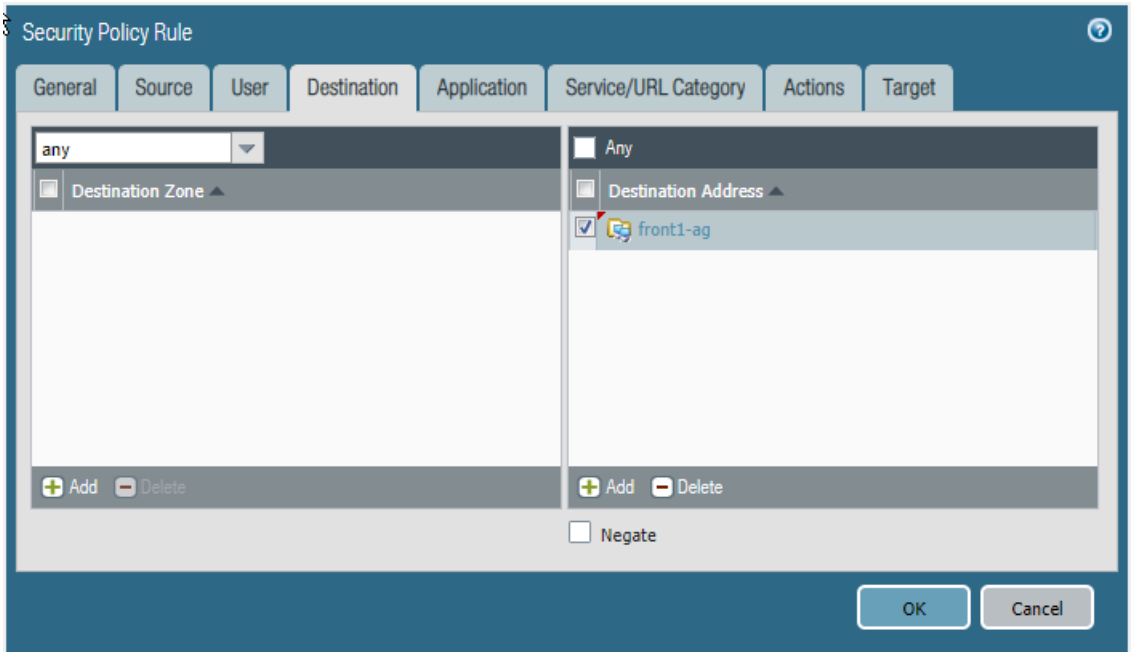

Deploy and Secure AKS Services

These steps outline how you can secure inbound and outbound traffic traversing to Kubernetes services using VM-Series firewall and the Azure Plugin for Panorama.

STEP 1 | In the application deployment environment, create a YAML file for the application or use a file that already exists. The following is a sample application YAML file:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: azure-vote-back
spec:
  replicas: 1
  selector:
    matchLabels:
      app: azure-vote-back
  template:
    metadata:
      labels:
```

 ©2024 Palo Alto Networks, Inc.

```
        app: azure-vote-back
    spec:
      containers:
      - name: azure-vote-back
        image: redis
        resources:
          requests:
            cpu: 100m
            memory: 128Mi
          limits:
            cpu: 250m
            memory: 256Mi
        ports:
        - containerPort: 6379
          name: redis
---
apiVersion: v1
kind: Service
metadata:
  name: azure-vote-back
  labels:
    service: backend
spec:
  ports:
  - port: 6379
  selector:
    app: azure-vote-back
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: azure-vote-front
spec:
  replicas: 5
  selector:
    matchLabels:
      app: azure-vote-front
  template:
    metadata:
      labels:
        app: azure-vote-front
    spec:
      containers:
      - name: azure-vote-front
        image: microsoft/azure-vote-front:v1
        resources:
          requests:
            cpu: 100m
            memory: 128Mi
          limits:
            cpu: 250m
            memory: 256Mi
        ports:
        - containerPort: 80
        env:
        - name: REDIS
```

 ©2024 Palo Alto Networks, Inc.

```
          value: "azure-vote-back"
---
apiVersion: v1
kind: Service
metadata:
  name: azure-vote-front
  labels:
    service: "azure-vote-front"
    type: "production"
    providesecurity: "yes"
    a: "value"
    b: "value"
    c: "value"
    tier: "stagingapp"
  annotations:
    service.beta.kubernetes.io/azure-load-balancer-internal: "true"
spec:
  type: LoadBalancer
  ports:
  - port: 80
  selector:
    app: azure-vote-front
```

STEP 2 | Edit your YAML file to label Kubernetes services.

Labels enable the corresponding tag-to-IP mapping to be created when you use the Panorama plugin for AKS to connect to the cluster. For example, in the above sample file look for the application labels in the service metadata. They are: **`azure-vote-back`** and **`azure-vote-front`**.

STEP 3 | In your AKS cluster, apply the YAML file.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | In Panorama, create an Address Group using a resource group tag.

1. On the **Objects** tab, select a device group from the **Device Group** list.
2. Select **Address Groups** and **Add** an Address Group.
   1. Specify a name, and select the **Dynamic** type.
   2. **Add** addresses. Adding spawns a window that lists detected addresses. Populating the list can take several minutes.
   3. You can choose one or more addresses for the Match Criteria. Select AND or OR for the criteria relationship.
   4. If you have many addresses, enter a string in the search box to filter the output, as shown in the following figure.
   5. In the address list, click the + to include the address in the address group match criteria.
   6. When the match criteria is complete, click **OK**.

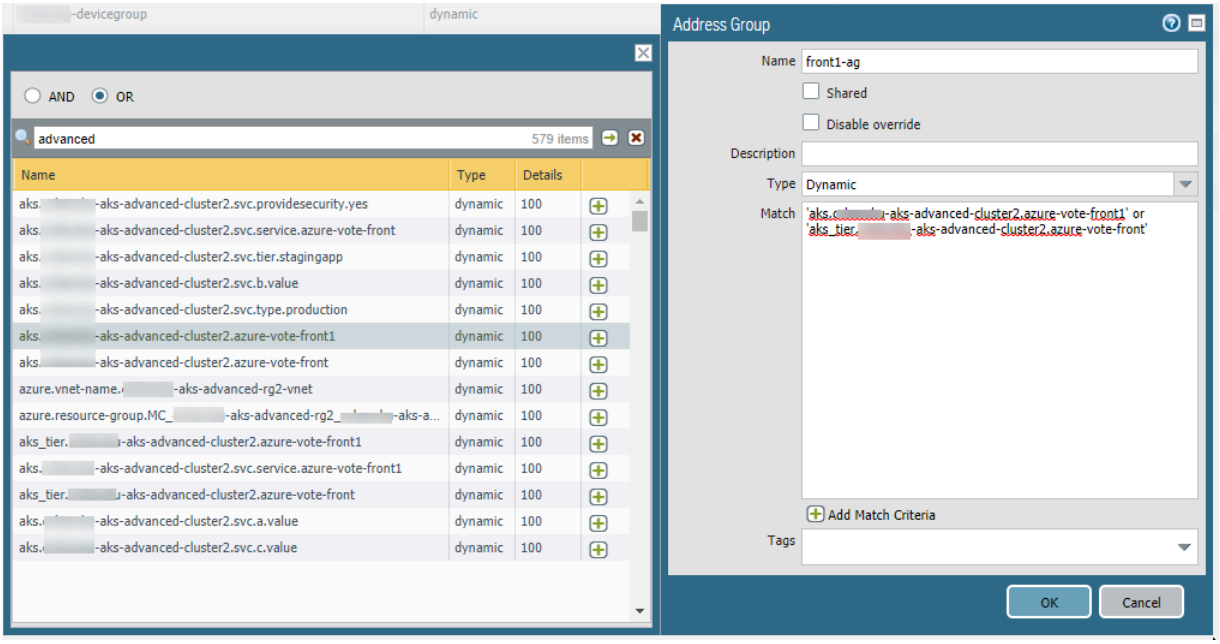

STEP 5 | Show Policy using the address group.

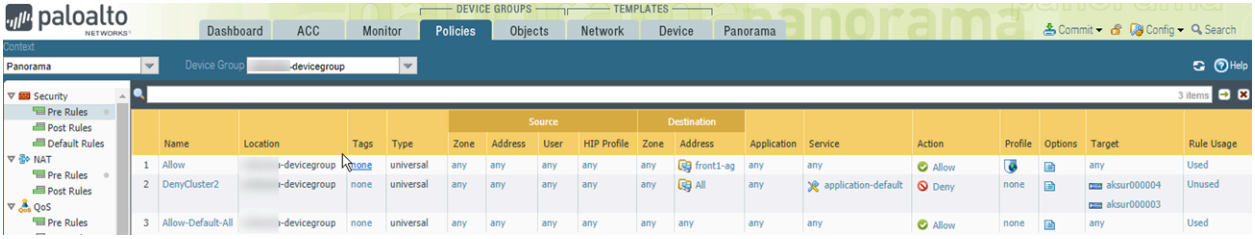

 ©2024 Palo Alto Networks, Inc.

STEP 6 | View secured AKS services.

In **Panorama** > **Azure** > **Deployments**, view your monitoring definition, and in the Action column select the **Protected Applications and Services** link.

The **Protected?** column summarizes the security status of your resource groups. It might take several minutes for the window to populate. If you have many resource groups, enter a string in the search box to filter the output.

This output is based on the Azure resource group configuration; it does not query the device group or template stack membership.

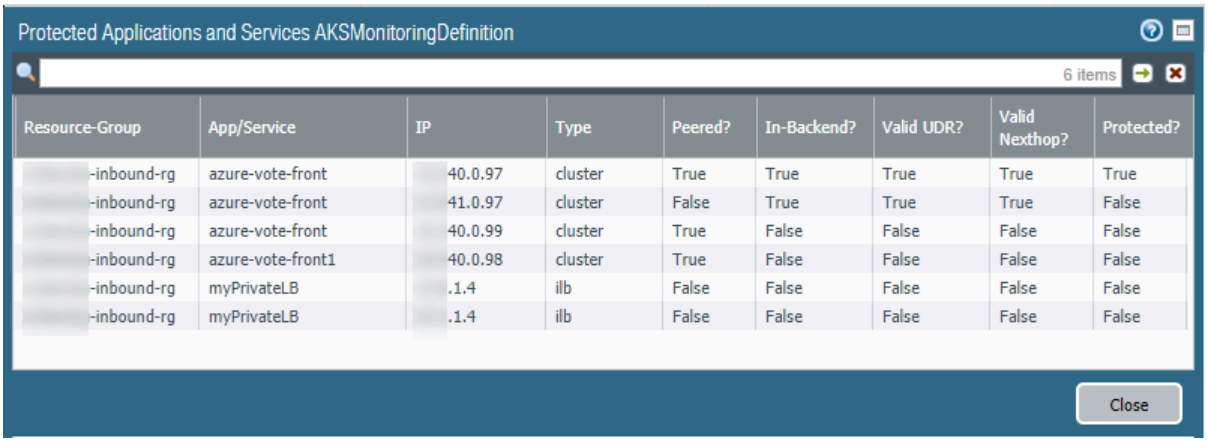

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on OpenStack

The VM-Series firewall for OpenStack allows you to provide secure application delivery along with network security, performance and visibility.

- VM-Series Firewall for OpenStack
- Components of the VM-Series for OpenStack Solution
- Heat Template for a Basic Gateway Deployment
- Heat Templates for Service Chaining and Service Scaling
- Install the VM-Series Firewall in a Basic Gateway Deployment
- Install the VM-Series Firewall with Service Chaining or Scaling

VM-Series Deployments in OpenStack

The Heat Orchestration templates provided by Palo Alto Networks allow you to deploy the VM-Series firewall individually, through service chaining, or dynamically with service scaling.

- Basic Gateway
- Service Chaining and Service Scaling

Basic Gateway

The VM-Series firewall for OpenStack allows you to deploy the VM-Series firewall on the KVM hypervisor running on a compute node in your OpenStack environment. This solution uses Heat Orchestration Templates and bootstrapping to deploy the VM-Series firewall and a Linux server. The VM-Series firewall protects the deployed Linux server by inspecting the traffic going in and out of the server. The sample bootstrap files allow the VM-Series firewall to boot with basic configuration for handling traffic.

These heat template files and the bootstrap files combine to create two virtual machines, the VM-Series firewall and Linux server, in a network configuration similar to that shown below.

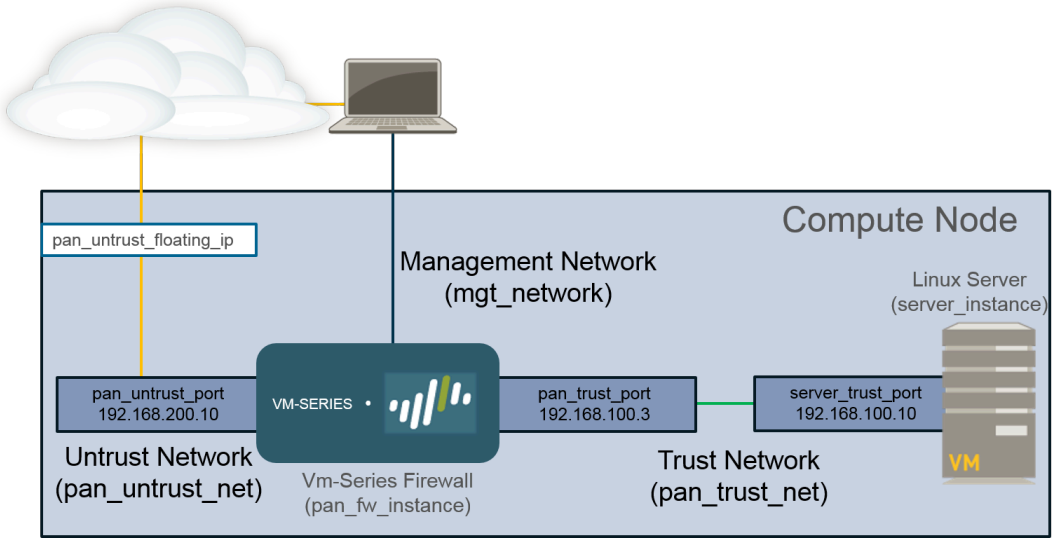

Service Chaining and Service Scaling

Deploying the VM-Series firewall through service chaining or service scaling is not supported on OpenStack Queens.

Service chaining is a Contrail feature that deploys a VM-Series firewall as a service instance in your OpenStack environment. A service chain is a set of service virtual machines, such as firewalls or load balancers, and each virtual machine in the service chain is a service instance. Service scaling allows you to dynamically deploy additional instances of the VM-Series firewall. Using CPU utilization or incoming bytes per second metrics gathered by Ceilometer, OpenStack deploys or shuts down additional instances of the VM-Series firewall to meet the current needs of your network.

 ©2024 Palo Alto Networks, Inc.

The VM-Series firewall in OpenStack solution leverages heat orchestration templates to configure and deploy the components required for service chaining and service scaling. The heat templates provided by Palo Alto networks create a service template, service instance, and service policy (to direct traffic to the VM-Series firewall) to deploy two Linux servers and the VM-Series firewall service instance between them.

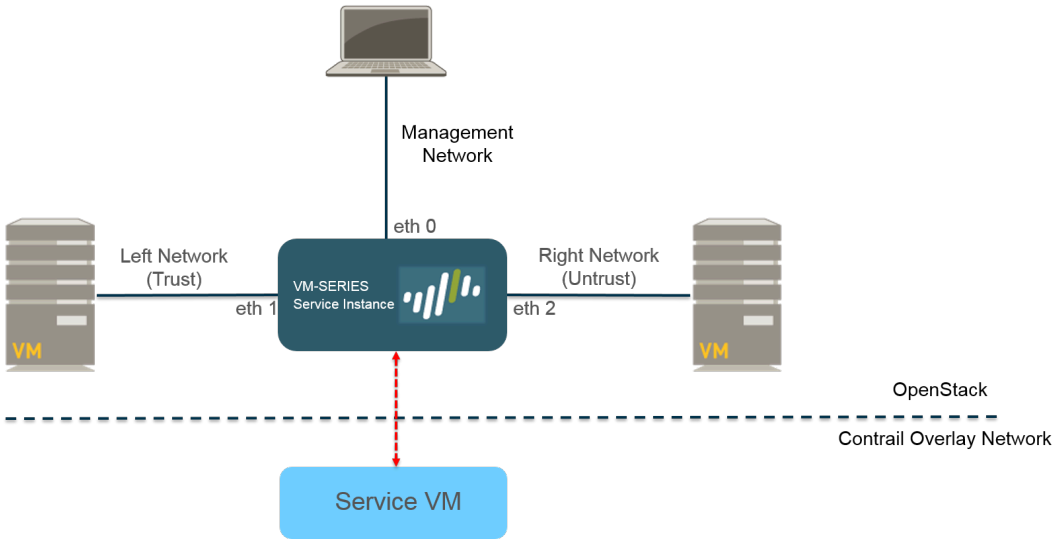

 ©2024 Palo Alto Networks, Inc.

Components of the VM-Series for OpenStack Solution

The VM-Series firewall in an OpenStack environment has been tested with the following components.

Component	Description
Software	See the Compatibility Matrix for details about supported software versions.
VM-Series Hardware Resources	See VM-Series System Requirements for the minimum hardware requirements for your VM-Series model. In OpenStack, flavors define the CPU, memory, and storage capacity of a compute instance. When setting up your Heat template, choose the compute flavor that meets or exceeds the hardware requirements for the VM-Series model.
Fuel Master	Fuel is a web UI-driven deployment and management tool for OpenStack.
OpenStack Controller	This node runs most of the shared OpenStack services, such API and scheduling. Additionally, the Horizon UI runs on this node.
OpenStack Compute	The compute node contains the virtual machines, including the VM-Series firewall, in the OpenStack deployment. The compute node that houses the VM-Series must meet the following criteria: • Instance type OS::Nova::Server • Allow configuration of at least three interfaces • Accept the VM-Series qcow2 image • Accept the compute flavor parameter *Install the OpenStack compute node on a bare-metal server because the VM-Series firewall does not support nested virtualization.*
Contrail Controller	The Contrail controller node is a software-defined networking controller used for management, control, and analytics for the virtualized network. It provides routing information to the compute and gateway nodes. Additionally, the Contrail controller provides the necessary support for service chaining and service scaling.
Contrail Gateway	The Contrail gateway node provides IP connectivity to external networks from virtual networks. MPLS over GRE tunnels from the

 ©2024 Palo Alto Networks, Inc.

Component	Description
	virtual machines terminate at the gateway node, where packets are decapsulated and sent to their destinations on IP networks.
Ceilometer (OpenStack Telemetry)	In the case of the VM-Series firewall for OpenStack, Ceilometer monitors CPU utilization for service scaling. When CPU utilization meets the defined thresholds, a new service instance of the VM-Series firewall is deployed or shut down.
Heat Orchestration Template Files	Palo Alto Networks provides a sample Heat template for deploying the VM-Series firewall. This template is made up of a main template and an environment template. These files instantiate one VM-Series instance with one management interface and two data interfaces. In a basic gateway deployment, the template instantiates a Linux server with one interface. The interface of the server attaches to the private network created by the template. In a service chaining or service scaling deployment, the templates instantiate two Linux servers with one server attached to each data interface of the firewall.
VM-Series Firewall Bootstrap Files	The VM-Series firewall bootstrap files consist of a init-cfg.txt file, bootstrap.xml file, and VM-Series auth codes. Along with the Heat template files, Palo Alto Networks provides a sample init-cfg.txt and bootstrap.xml files. You must provide your own auth codes to license your VM-Series firewall and activate any subscriptions. See Bootstrap the VM-Series Firewall for more information about VM-Series bootstrap files.

 ©2024 Palo Alto Networks, Inc.

Heat Template for a Basic Gateway Deployment

The heat template file includes the following four files to help you launch the VM-Series firewall on KVM in OpenStack. All four files are required to deploy the VM-Series firewall and Linux server.

- **pan_basic_gw.yaml**—Defines the resources created to support the VM-Series firewall and Linux server on the compute node, such as interfaces and IP addresses.
- **pan_basic_gw_env.yaml**—Defines the environment that the VM-Series firewall and Linux server exist in. Many parameters in the pan_basic_gw.yaml file reference the parameters defined in this file, such as flavor for the VM-Series and the Linux server.
- **init-cfg.txt**—Includes the operational command to enable DHCP on the firewall management interface.
- **bootstrap.xml**—Provides basic configuration for the VM-Series firewall. The bootstrap.xml file configures the data interfaces and IP addresses. These values must match the corresponding values in the pan_basic_gw.yaml file.

  Additionally, the bootstrap.xml file includes a NAT rule called untrust2trust. This rule translate the trust port on the server to the untrust port of the VM-Series firewall.

You have two options for passing bootstrapping files to OpenStack—file injection (personality files) or user data.

File injection is no longer supported beginning with OpenStack Queens; you must use user data instead.

The table below describes resources that the pan_basic_gw.yaml template file creates and provides the default value, if applicable.

Resource	Description
pan_fw_instance	VM-Series firewall with a management interface and two data interfaces.
server_instance	A Linux server with a single interface.
pan_trust_net	A connection to the internal network to which the trust interface of the firewall and trust interface of the server are attached.
pan_trust_subnet	Subnet attached to the trust interface on the firewall (pan_trust_net) and has a CIDR value of 192.168.100.0/24.
pan_untrust_net	Untrust network to which the untrust port of the firewall is attached.
pan_untrust_subnet	Subnet attached to the untrust interface of the firewall (pan_untrust_net) and has a CIDR value of 192.168.200.0/24.

 ©2024 Palo Alto Networks, Inc.

Resource	Description
allow_ssh_https_icmp_secgroup	Security group that allows TCP on ports 22 and 443 and ICMP traffic.
pan_untrust_port	The untrust port of the VM-Series firewall deployed in Layer 3 mode. The Heat template provides a default IP address of 192.168.200.10 to this port. If you change this IP address in the heat template, you must change the IP address in the bootstrap.xml file.
pan_untrust_floating_ip	A floating IP address assigned from the public_network.
pan_untrust_floating_ip_assoc	This associates the pan_untrust_floating_ip to the pan_untrust_port.
pan_trust_port	The trust port of the VM-Series firewall Layer 3 mode.
server_trust_port	The trust port of the Linux server Layer 3 mode. The Heat template provides a default IP address of 192.168.100.10 to this port. If you change this IP address in the heat template, you must change the IP address in the bootstrap.xml file.

The pan_basic_gw.yaml file references the pan_basic_gw_env.yaml for many of the values needed to create the resources need to deploy the VM-Series firewall and Linux server. The heat template environment file contains the following parameters.

Parameter	Description
mgmt_network	The VM-Series firewall management interface attaches to the network specified in this parameter. The template does not create the management network; you must create this before deploying the heat templates. The default value is mgmt_ext_net.
public_network	Addresses that the OpenStack cluster and the virtual machines in the cluster use to communicate with the external or public network. The public network provides virtual IP addresses for public endpoints, which are used to connect to OpenStack services APIs. The template does not create the public network; you must create this before deploying the heat templates. The default value is public_net.
pan_image	This parameter specifies the VM-Series base image used by the Heat template when deploying the VM-Series firewall. The default value is pa-vm-7.1.4.

 ©2024 Palo Alto Networks, Inc.

Parameter	Description
pan_flavor	This parameter defines the hardware resources allocated to the VM-Series firewall. The default value is m1.medium. This value meets the VM-Series on KVM System Requirements described in the Set Up the VM-Series Firewall on KVM chapter.
server_image	This parameter tells the Heat template which image to use for the Linux server. The default value is Ubuntu-14.04.
server_flavor	This parameter defines the hardware resources allocated to the Linux server. The default value is m1.small.
server_key	The server key is used for accessing the Linux server through ssh. The default value is server_key. You can change this value by entering a new server key in the environment file.

 ©2024 Palo Alto Networks, Inc.

Heat Templates for Service Chaining and Service Scaling

Deploying the VM-Series firewall through service chaining or service scaling is not supported on OpenStack Queens.

The heat template environment file defines the parameters specific to the VM-Series firewall instance deployed through service chaining or service scaling. The parameters defined in the environment file are divided into sections described below. There are two versions of the heat templates for service chaining—vwire and L3— and one for service scaling.

Service chaining requires the heat template files and two bootstrap files to launch the VM-Series firewall service instance and two Linux servers in the left and right networks.

- **Template files**—This template defines the resources created to support the VM-Series firewall and two Linux servers, such as interfaces and IP addresses.
  - service_chaining_template_vm.yaml for vwire deployments.
  - service_chaining_template_L3.yaml for L3 deployments.
  - service_scaling_template.yaml for service scaling deployments.
- **Environment file**—This environment file defines the environment that the VM-Series firewall and Linux servers exist in. Many parameters in the template reference the parameters defined in this file, such as flavor for the VM-Series and the names of the Linux servers.
  - service_chaining_env_vm.yaml for vwire deployments.
  - service_chaining_env_L3.yaml for L3 deployments.
  - service_scaling_env.yaml for service scaling deployments.
- **service_instance.yaml**—(Service Scaling only) This is a nested heat template that is reference by Service_Scaling_template.yaml to deploy the service instance. It provides the necessary information to deploy service instances for scaling events.
- **init-cfg.txt**—Provides the minimum information required to bootstrap a VM-Series firewall. The init-cfg.txt provided only includes the operational command to enable DHCP on the firewall management interface.
- **<file_name>_bootstrap.xml**—Provides basic configuration for the VM-Series firewall. The bootstrap.xml file configures the data interfaces. These values must match the corresponding values in the heat templates files.

For more information about the init-cfg.txt and bootstrap.xml files, see Bootstrap Configuration Files.

The following tables describe the parameters of the environment file.

- Virtual Network
- Virtual Machine
- Service Template
- Service Instance
- IPAM
- Service Policy

 ©2024 Palo Alto Networks, Inc.

- Alarm

Virtual Network

The virtual network configuration paramaters in the heat template environment file define the virtual network that connects the VM-Series firewall and the two Linux servers deployed by the heat template.

Virtual Network (VN Config)	
management_network	The VM-Series firewall management interface attaches to the network specified in this parameter.
left_vn or left_network	Name of the left virtual network.
right_vn or right_network	Name of the right virtual network.
left_vn_fqdn	Fully qualified domain name of the left virtual network.
right_vn_fqdn	Fully qualified domain name of the right virtual network
route_target	Edit this value so route target configuration matches that of your external gateway.

Virtual Machine

The virtual machine parameters define the left and right Linux servers. The name of the port tuple is defined here and referenced by the heat template. In Contrail, a port tuple is an ordered set of virtual network interfaces connected to the same virtual machine. With a port tuple, you can create ports and pass that information when creating a service instance. The heat template creates the left, right, and management ports and adds them to the port tuple. The port tuple is then linked to the service instance. When you launch the service instance using the heat templates, the port tuple maps the service virtual machine to the virtual machine deployed in OpenStack.

Virtual Machine (VM Config)	
flavor	The flavor of the left and right virtual machines. The default value is m1.small.
left_vm_image or right_vm_image or image	The name of the software image for the left and right virtual machines. Change this value to match the file name of the image you uploaded. The default is TestVM, which is a default image provided by OpenStack.

 ©2024 Palo Alto Networks, Inc.

Virtual Machine (VM Config)	
svm_name	The name applied to the VM-Series firewall.
left_vm_name and right_vm_name	The name of the left and right virtual machines.
port_tuple_name	The name of the port tuple used by the two Linux servers and the VM-Series firewall.
server_key	The server key is used for accessing virtual machines through SSH. The default value is server_key. You can change this value by entering a new server key in the environment file.

Service Template

The service template defines the parameters of the service instance, such as the software image, virtual machine flavor, service type, and interfaces. Service templates are configured within the scope of a domain and can be used on all projects within the specified domain.

Service Template (ST Config)	
S_Tmp_name	The name of the service template.
S_Tmp_version	The service template version. The default value is 2. Do not change this parameter because service template version 2 is required to support port tuples.
S_Tmp_service_mode	Service mode is the network mode used by the VM-Series firewall service instance. For the L3 network template, the default value is in-network. For the virtual wire template, the default value is transparent.
S_Tmp_service_type	The type of service being deployed by the template. The default value is firewall and should not be changed when deploying the VM-Series firewall.
S_Tmp_image_name	This parameter specifies the VM-Series base image used by the Heat template when deploying the VM-Series firewall. Edit this parameter to match the name of the VM-Series firewall image uploaded to your OpenStack environment.
S_Tmp_flavor	This parameter defines the hardware resources allocated to the VM-Series firewall. The default value is m1.large.
S_Tmp_interface_type_mgmt S_Tmp_interface_type_left S_Tmp_interface_type_right	These parameters define the interface type for management, left, and right interfaces.

 ©2024 Palo Alto Networks, Inc.

Service Template (ST Config)	
domain	The domain where this service template is tied to. The default value is default-domain.

Service Instance

The service instance portion of the heat template environment file provides the name of the individual instance deployed by the heat template and service template.

Service Instance (SI Config)	
S_Ins_name	The service instance name. This is the name of the VM-Series firewall instance in Contrail.
S_Ins_fq_name	The fully qualified name of the service instance.

IPAM

IP address management (IPAM) provides the IP address information for the interfaces of the service instance. Changes these parameters to best suit your environment.

IPAM (IPAM Config)	
NetIPam_ip_prefix_mgmt	The IP prefix of the management interface on the VM-Series firewall. The default value is 172.2.0.0.
NetIPam_ip_prefix_len_mgmt	The IP prefix length of the management interface on the VM-Series firewall. The default value is /24.
NetIPam_ip_prefix_left	The IP prefix of the left interface on the VM-Series firewall. The default value is 10.10.1.0.
NetIPam_ip_prefix_len_left	The IP prefix length of the left interface on the VM-Series firewall. The default value is /24.
NetIPam_ip_prefix_right	The IP prefix of the right interface on the VM-Series firewall. The default value is 10.10.2.0.
NetIPam_ip_prefix_len_right	The IP prefix length of the right interface on the VM-Series firewall. The default value is /24.
NetIPam_addr_from_start_flag	This parameter determines how IP addresses are assigned to VMs on the subnets described above. If true, any new VM takes the next available IP address. If false, any new VM is assigned an IP address at random. The default value is true.

 ©2024 Palo Alto Networks, Inc.

Service Policy

The service policy defines the traffic redirection rules and policy that point traffic passing between the left and right virtual machines to the VM-Series firewall service instance.

Service Policy (Policy Config)	
policy_name	The name of the service policy in Contrail that redirects traffic through the VM-Series firewall. For the L3 template, the default value is PAN_SVM_policy-L3. For the virtual wire template, the default value is PAN_SVM_policy-vw.
policy_fq_name	The fully qualified name of the service policy.
simple_action	The default action Contrail applies to traffic going to the VM-Series firewall service instance. The default value is pass because the VM-Series firewall will apply its own security policy to the traffic.
protocol	The protocols allowed by Contrail to pass to the VM-Series firewall. The default value is any.
src_port_end and src_port_start	Use this parameter to specify source port(s) that should be associated with the policy rule. You can enter a single port, a list of ports separated with commas, or a range of ports in the form of <port>-<port>. The default value is -1 in the provided heat templates; meaning any source port.
direction	This parameter defines the direction of traffic that is allowed by Contrail to pass to the VM-Series firewall. The default value is <> or bidirectional traffic.
dst_port_end and dst_port_start	Use this parameter to specify destination port(s) that should be associated with the policy rule. You can enter a single port, a list of ports separated with commas, or a range of ports in the form of <port>-<port>. The default value is -1 in the provided heat templates; meaning any destination port.

Alarm

The alarm parameters are used in service scaling and are not included in the service chaining environment files. These parameters define the thresholds used by Contrail to determine when scaling should take place. This set of parameters is only used the service scaling heat template.

The default time configured under the cooldown parameters is intended to allow the firewall enough time to boot up. If you change the cooldown values, leave sufficient time for each new firewall instance to boot up.

 ©2024 Palo Alto Networks, Inc.

Alarm	
meter_name	The metric monitored by Ceilometer and used by contrail to determine when an additional VM-Series firewall should be deployed or brought down. The heat template uses CPU utilization or bytes per second as metrics for service scaling.
cooldown_initial	The amount time Contrail waits before launching a additional service instance after the initial service instance is launched. The default is 1200 seconds.
cooldown_scaleup	The amount of time Contrail waits between launching additional service instance after the first scale up service instance launch. The default is 1200 seconds.
cooldown_scaledown	The amount of time Contrail waits between shutting down additional service instances after the first scale up service instance shut down. The default is 1200 seconds.
period_high	The interval during which the average CPU load is calculated as high before triggering an alarm. The default value is 300 seconds.
period_low	The interval during which the average CPU load is calculated as low before triggering an alarm. The default value is 300 seconds.
threshold_high	The value of CPU utilization in percentage or bytes per second that Contrail references before launching a scale up event. The default is 40% CPU utilization or 2800 bytes per second.
threshold_low	The value of CPU utilization in percentage or bytes per second that Contrail references before launching a scale down event. The default is 20% CPU utilization or 12000 bytes per second.

 ©2024 Palo Alto Networks, Inc.

Install the VM-Series Firewall in a Basic Gateway Deployment

Complete the following steps to prepare the heat templates, bootstrap files, and software images needed to deploy the VM-Series firewall in OpenStack. After preparing the files, deploy the VM-Series firewall and Linux server.

STEP 1 | Download the Heat template and bootstrap files.

Download the Heat template package from the GitHub repository.

STEP 2 | Download the VM-Series base image.

1. Login in to the Palo Alto Networks Customer Support Portal.
2. Select **Software Updates** and choose **PAN-OS for VM-Series KVM Base Images** from the **Filter By** drop-down.
3. Download the VM-Series for KVM**qcow2** file.

STEP 3 | Download Ubuntu 14.04 and upload the image to the OpenStack controller.

The Heat template needs an Ubuntu image for launching the Linux server.

1. Download Ubuntu 14.04.
2. Log in to the Horizon UI.
3. Select **Project** > **Compute** > **Images** > **Create Image**.
4. **Name** the image Ubuntu 14.04 to match the parameter in the pan_basic_gw_env.yaml file.
5. Set Image Source to **Image File**.
6. Click **Choose File** and navigate to your Ubuntu image file.
7. Set the Format to match the file format of your Ubuntu image.
8. Click **Create Image**.

STEP 4 | Upload the VM-Series for KVM base image to the OpenStack controller.

1. Log in to the Horizon UI.
2. Select **Project** > **Compute** > **Images** > **Create Image**.
3. **Name** the image to match the image name in your Heat template.
4. Set Image Source to **Image File**.
5. Click **Choose File** and navigate to your VM-Series image file.
6. Set the Format to **QCOW2-QEMU Emulator**.
7. Click **Create Image**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Upload the bootstrap files. You have two options for passing bootstrapping files to OpenStack—file injection (personality files) or user data. To pass the bootstrap files using user-data, you must place the files in a tar ball (.tgz file) and encode that tar ball with base64.

> *File injection is no longer supported beginning with OpenStack Queens; you must use user data instead.*

- For file injection, upload the init-cfg.txt, bootstrap.xml, and your VM-Series auth codes to your OpenStack controller or a web server that the OpenStack controller can access.
- If using the **--user-data** method to pass the bootstrap package to the config-drive, you can use the following command to create the tar ball and encode the tar ball (.tgz file) with base64:

```
tar -cvzf <file-name>.tgz config/
license software content
base64 -i <in-file> -o <outfile>
```

STEP 6 | Edit the pan_basic_gw.yaml template to point to the bootstrap files and auth codes.

- If you are using personality files, specify the file path or web server address to the location of your files under personality. Uncomment whichever lines you are not using.

```
  pan_fw_instance:
    type: OS::Nova::Server
    properties:
      image: { get_param: pan_image }
      flavor: { get_param: pan_flavor }
      networks:
        - network: { get_param: mgmt_network }
        - port: { get_resource: pan_untrust_port }
        - port: { get_resource: pan_trust_port }
      user_data_format: RAW
      config_drive: true
      personality:
        /config/init-cfg.txt: {get_file: "/opt/pan_bs/init-
cfg.txt"}
#        /config/init-cfg.txt: { get_file: "http://
web_server_name_ip/pan_bs/init-cfg.txt" }
        /config/bootstrap.xml: {get_file: "/opt/pan_bs/
bootstrap.xml"}
#        /config/bootstrap.xml: { get_file: "http://
web_server_name_ip/pan_bs/bootstrap.xml" }
        /license/authcodes: {get_file: "/opt/pan_bs/authcodes"}
#        /license/authcodes: {get_file: "http://
web_server_name_ip/pan_bs/authcodes"}
```

- If you are using user-data, specify the file path or web server address to the location of your files under user_data. If you have more than one

```
pan_fw_instance:
    type: OS::Nova::Server
    properties:
```

 ©2024 Palo Alto Networks, Inc.

```
      image: { get_param: pan_image }
      flavor: { get_param: pan_flavor }
      networks:
        - port: { get_resource: mgmt_port }
        - port: { get_resource: pan_untrust_port }
        - port: { get_resource: pan_trust_port }
      user_data_format: RAW
      config_drive: true
      user_data:
#        get_file: http://10.0.2.100/pub/repository/panos/images/
openstack/userdata/boot.tgz
       get_file: /home/stack/newhot/bootfiles.tgz
```

STEP 7 | Edit the pan_basic_gw_env.yaml template environment file to suit your environment. Make sure that the management and public network values match those that you created in your OpenStack environment. Set the pan_image to match the name you assigned to the VM-Series base image file. You can also change your server key here.

```
root@node-2:~# cat basic_gateway/pan_basic_gw_env.yaml
parameters:
  mgmt_network: mgmt_ext_net
  public_network: public_net
  pan_image: pa-vm-image
  pan_flavor: m1.medium
  server_image: Ubuntu-14.04
  server_flavor: m1.small
  server_key: server_key
```

STEP 8 | Deploy the Heat template.

1. Execute the command **`source openrc`**
2. Execute the command **`heat stack-create <stack-name> -f <template> -e ./<env-template>`**

```
root@node-2:~# heat stack-create stack1 -f pan_basic_gw.yaml -e pan_basic_gw_env.yaml
+--------------------------------------+------------+--------------------+---------------------+--------------+
| id                                   | stack_name | stack_status       | creation_time       | updated_time |
+--------------------------------------+------------+--------------------+---------------------+--------------+
| ebe40f9d-2781-4bb2-b246-f15c761f9045 |   stack1   | CREATE_IN_PROGRESS | 2017-01-25T13:36:59 | None         |
+--------------------------------------+------------+--------------------+---------------------+--------------+
```

STEP 9 | Verify that your VM-Series firewall is deployed successfully.

You can use the following commands to check the creation status of the stack.

- Check the stack status with **`heat stack-list`**
- View a detailed list of events that occurred during stack creation with **`heat event-list`**
- View details about your stack with **`heat stack-show`**

STEP 10 | Verify that the VM-Series firewall is bidirectionally inspecting traffic accessing the Linux server.

1. Log in to the firewall.
2. Select **Monitor** > **Logs** > **Traffic** to view the SSH session.

 ©2024 Palo Alto Networks, Inc.

Install the VM-Series Firewall with Service Chaining or Scaling

Complete the following steps to prepare the heat templates, bootstrap files, and software images needed to deploy the VM-Series firewall. After preparing the files, deploy the VM-Series firewall service and two Linux servers.

Deploying the VM-Series firewall through service chaining or service scaling is not supported on OpenStack Queens.

STEP 1 | Download the Heat template and bootstrap files.

Download the Heat template package from the GitHub repository.

STEP 2 | Download the VM-Series base image.

1. Login in to the Palo Alto Networks Customer Support Portal.
2. Select **Software Updates** and choose **PAN-OS for VM-Series KVM Base Images** from the **Filter By** drop-down.
3. Download the VM-Series for KVM**qcow2** file.

STEP 3 | Download Ubuntu 14.04 and upload the image to the OpenStack controller.

For service chaining, you can use the default image provided by OpenStack called TestVM. Skip this step when using TestVM. An Ubuntu image is required for service scaling.

1. Download Ubuntu 14.04.
2. Log in to the Horizon UI.
3. Select **Project** > **Compute** > **Images** > **Create Image**.
4. **Name** the image Ubuntu 14.04 to match the parameter in the pan_basic_gw_env.yaml file.
5. Set Image Source to **Image File**.
6. Click **Choose File** and navigate to your Ubuntu image file.
7. Set the Format to match the file format of your Ubuntu image.
8. Click **Create Image**.

A server key is required when using an Ubuntu image. Ensure that the server key is added to the environment file.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Upload the VM-Series for KVM base image to the OpenStack controller.

1. Log in to the Horizon UI.
2. Select **Project** > **Compute** > **Images** > **Create Image**.
3. **Name** the image to match the image name in your Heat template.
4. Set Image Source to **Image File**.
5. Click **Choose File** and navigate to your VM-Series image file.
6. Set the Format to **QCOW2-QEMU Emulator**.
7. Click **Create Image**.

STEP 5 | Upload the bootstrap files. The files must be uploaded to the folder structure described here. The heat template uses this folder structure to locate the bootstrap files.

1. Log in to your OpenStack controller.
2. Create the following folder structure:

   /root/bootstrap/config/

   /root/bootstrap/license/
3. Using SCP or FTP, add the init-cfg.txt and bootstrap.xml files to the config folder and add your VM-Series auth codes to the license folder.

STEP 6 | Edit the template environment file to suit your environment. Verify that the image names in the environment file match the names you gave the files when you uploaded them.

```
parameters:
# VN config
  management_network: 'mgmt_net'
  left_vn: 'left_net'
  right_vn: 'right_net'
  left_vn_fqdn: 'default-domain:admin:left_net'
  right_vn_fqdn: 'default-domain:admin:right_net'
  route_target: "target:64512:20000"
# VM config
  flavor: 'm1.small'
  left_vm_image: 'TestVM'
  right_vm_image: 'TestVM'
  svm_name: 'PAN_SVM_L3'
  left_vm_name: 'Left_VM_L3'
  right_vm_name: 'Right_VM_L3'
  port_tuple_name: 'port_tuple_L3'
#ST Config
  S_Tmp_name: PAN_SVM_template_L3
  S_Tmp_version: 2
  S_Tmp_service_mode: 'in-network'
  S_Tmp_service_type: 'firewall'
  S_Tmp_image_name: 'PA-VM-8.0.0'
  S_Tmp_flavor: 'm1.large'
  S_Tmp_interface_type_mgmt: 'management'
  S_Tmp_interface_type_left: 'left'
  S_Tmp_interface_type_right: 'right'
  domain: 'default-domain'
# SI Config
```

 ©2024 Palo Alto Networks, Inc.

```
  S_Ins_name: PAN_SVM_Instance_L3
  S_Ins_fq_name: 'default-domain:admin:PAN_SVM_Instance_L3'
#IPAM Config
  NetIPam_ip_prefix_mgmt: '172.2.0.0'
  NetIPam_ip_prefix_len_mgmt: 24
  NetIPam_ip_prefix_left: '10.10.1.0'
  NetIPam_ip_prefix_len_left: 24
  NetIPam_ip_prefix_right: '10.10.2.0'
  NetIPam_ip_prefix_len_right: 24
  NetIPam_addr_from_start_true: true
#Policy Config
  policy_name: 'PAN_SVM_policy-L3'
  policy_fq_name: 'default-domain:admin:PAN_SVM_policy-L3'
  simple_action: 'pass'
  protocol: 'any'
  src_port_end: -1
  src_port_start: -1
  direction: '< >'
  dst_port_end: -1
  dst_port_start: -1
```

STEP 7 | Edit the template files to point to the bootstrap files and auth codes. Under Personality, specify the file path to the location of your files. Uncomment whichever lines you are not using.

```
Pan_Svm_instance:
    type: OS::Nova::Server
    depends_on: [ mgmt_InstanceIp, left_InstanceIp,
 right_InstanceIp ]
    properties:
      name: {get_param: svm_name }
      image: { get_param:  S_Tmp_image_name }
      flavor: { get_param: S_Tmp_flavor }
      networks:
        - port: { get_resource: mgmt_VirtualMachineInterface }
        - port: { get_resource: left_VirtualMachineInterface }
        - port: { get_resource: right_VirtualMachineInterface }
      user_data_format: RAW
      config_drive: true
      personality:
        /config/init-cfg.txt: {get_file: "/root/bootstrap/config/
init-cfg.txt"}
#        /config/init-cfg.txt: { get_file: "http://10.4.1.21/
op_test/config/init-cfg.txt" }
        /config/bootstrap.xml: {get_file: "/root/bootstrap/config/
Service_Chaining_bootstrap_L3.xml"}
#        /config/bootstrap.xml: { get_file: "http://10.4.1.21/
op_test/config/Service_Chaining_bootstrap_L3.xml" }
#        /license/authcodes: {get_file: "/root/bootstrap/license/
authcodes"}
#        /license/authcodes: {get_file: "http://10.4.1.21/op_test/
license/authcodes"}
```

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Upload the heat template files.

1. Log in to your OpenStack Controller.
2. Use SCP or FTP to add the heat template file and environment file.

STEP 9 | Deploy the Heat template.

1. Execute the command **`source openrc`**
2. Execute the command **`heat stack-create <stack-name> -f <template> -e ./<env-template>`**

STEP 10 | Verify that your VM-Series firewall is deployed successfully.

You can use the following commands to check the creation status of the stack.

- Check the stack status with **`heat stack-list`**
- View a detailed list of events that occurred during stack creation with **`heat event-list`**
- View details about your stack with **`heat stack-show`**

STEP 11 | Verify that the VM-Series firewall is bidirectionally inspecting traffic between the Linux servers.

1. Log in to the firewall.
2. Select **Monitor** > **Logs** > **Traffic** to view the SSH session.

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on Google Cloud Platform

You can deploy a VM-Series firewall on a Google Compute Engine instance on the Google Cloud Platform.

- Supported Deployments on Google Cloud Platform
- Prepare to Set Up the VM-Series Firewall on Google Public Cloud
- Deploy the VM-Series Firewall on Google Cloud Platform
- VM Monitoring with the Panorama Plugin for GCP
- Auto Scaling the VM-Series Firewall on Google Cloud Platform

About the VM-Series Firewall on Google Cloud Platform

VM-Series firewalls bring next-generation firewall features to the Google® Cloud Platform (GCP™).

To maximize performance, VM-Series firewalls on GCP support the Data Plane Development Kit (DPDK) libraries, which provide fast packet processing and improve network performance based on specific combinations of VM-Series firewall licenses and Google Cloud Platform virtual machine (VM) sizes.

- Google Cloud Platform and the VM-Series Firewall
- Minimum System Requirements for the VM-Series Firewall

Google Cloud Platform and the VM-Series Firewall

The VM-Series firewall integration with Google Cloud Platform (GCP) allows you to deploy the VM-Series firewall as a virtual machine (VM) running on a Google Compute Engine instance. This process is simplified when you Deploy the VM-Series Firewall from Google Cloud Platform Marketplace.

After you deploy the VM-Series firewall, you can configure the following optional services:

- Enable Google Stackdriver Monitoring on the VM Series Firewall—From the firewall, push PAN-OS metrics to the Google Stackdriver service.
- Enable VM Monitoring to Track VM Changes on Google Cloud Platform—Set up a VM information source that monitors the specific GCP zone containing your instances. The monitored VM metadata can include predefined GCP properties (such as the project ID) and user-defined properties (such as labels and network tags).

Minimum System Requirements for the VM-Series Firewall

You must choose a VM-Series Firewall License for Public Clouds and a license method: bring-your-own-license (BYOL) or pay-as-you-go (PAYG). To deploy a VM-Series firewall on a Google Compute Engine instance, you must choose a machine type that supports the VM-Series System Requirements for your license.

A single Google Compute Engine instance supports up to eight network interfaces. If you want to configure eight interfaces, choose n1-standard-8 or a larger machine type.

Capacity	BYOL	Bundles 1 and 2	
		PAYG	Marketplace
VM-100 Firewall	✓		
VM-200 Firewall	✓		
VM-300 Firewall	✓	✓	✓

 ©2024 Palo Alto Networks, Inc.

Capacity	BYOL	Bundles 1 and 2	
		PAYG	Marketplace
VM-1000-HV Firewall	✓		
VM-500 Firewall	✓		
VM-700 Firewall	✓		

The VM-Series firewall supports the predefined standard machine types listed below. You can choose a higher performing machine type or you can create your own custom machine type if the resource requirements are compatible with your VM-Series firewall license.

- n1-standard-4
- n1-standard-8
- n1-standard-16
- n2-standard-4
- n2-standard-8
- n2-standard-16
- n2-standard-32

 ©2024 Palo Alto Networks, Inc.

Supported Deployments on Google Cloud Platform

You can deploy the VM-Series firewall on a Google® Compute Engine instance in a network in your virtual private cloud (VPC). The deployment types are:

- Internet Gateway
- Segmentation Gateway
- Hybrid IPSec VPN

Internet Gateway

The VM-Series firewall secures North/South traffic to and from the internet to protect applications from known and unknown threats. A Google project can have up to five VPC networks. For a typical example of an internet gateway, refer to the Google configuration examples.

In public cloud environments, it is a common practice to use a scale-out architecture (see the figure below) rather than larger, higher performing VMs. This architecture (sometimes called a *sandwich* deployment) avoids a single point of failure and enables you to add or remove firewalls as needed.

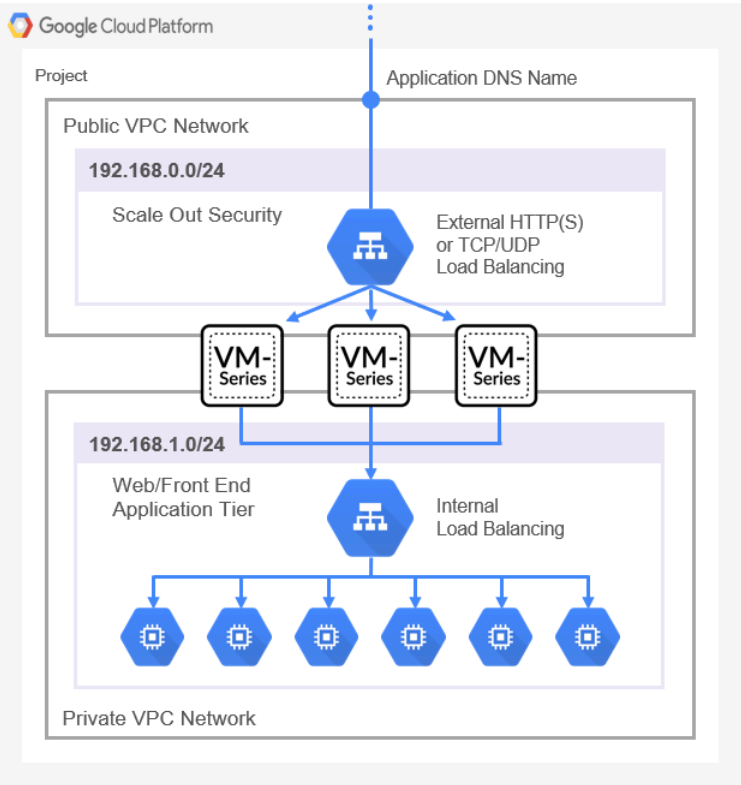

Segmentation Gateway

A segmentation gateway secures East/West traffic between virtual private clouds (VPCs) to ensure data protection compliance and application access. The following figure shows a firewall securing both North/South and East/West traffic.

 ©2024 Palo Alto Networks, Inc.

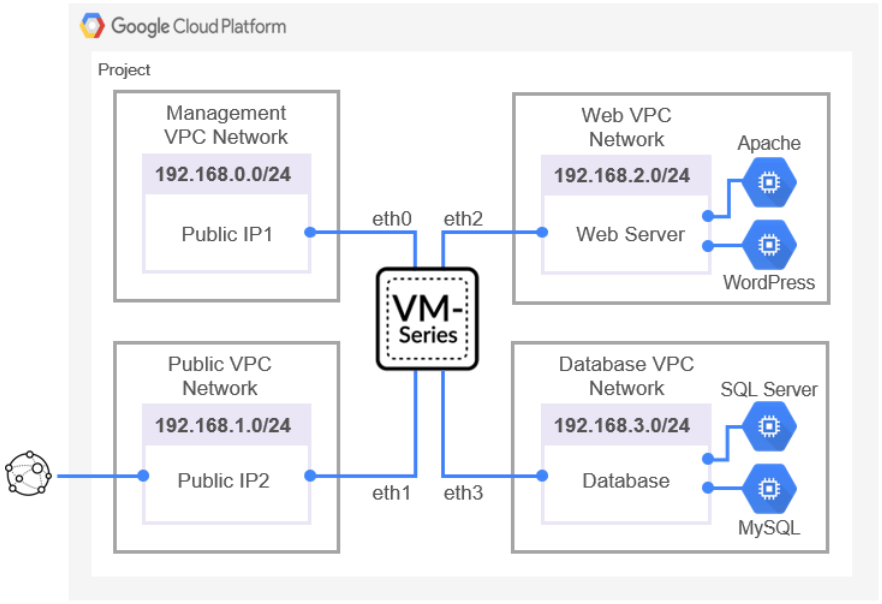

Hybrid IPSec VPN

The VM-Series firewall serves as an IPSec VPN termination point, which enables secure communications to and from applications hosted on Google Cloud Platform (GCP).

The deployment in the figure below shows a site-to-site VPN from an on-premises network to a VM-Series firewall deployed on GCP and an IPSec connection from an on-premises network to a Google Cloud VPN gateway.

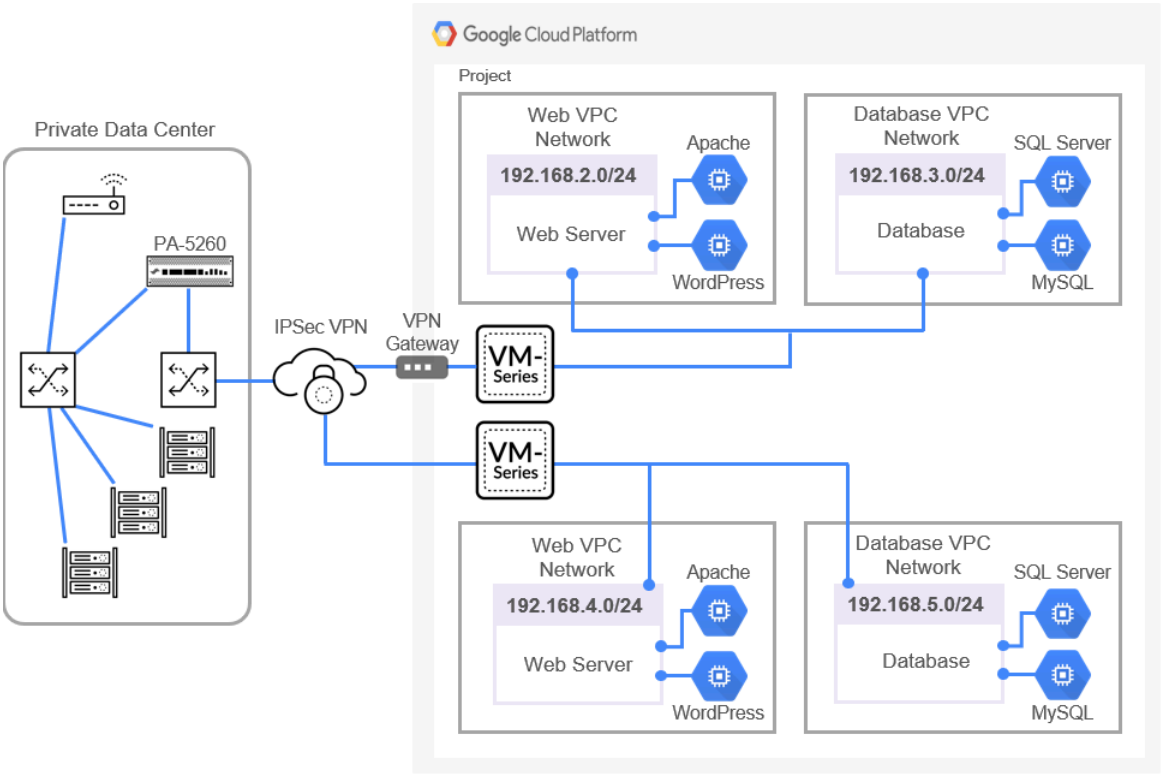

 ©2024 Palo Alto Networks, Inc.

Prepare to Set Up VM-Series Firewalls on Google Public Cloud

The process to Deploy the VM-Series Firewall from Google Cloud Platform Marketplace requires preparation tasks.

If you are deploying using the Google Marketplace, you must create your project networks and subnetworks, and plan networks and IP address assignments for the VM-Series firewall interfaces in advance. During the deployment, you must choose from existing networks and subnetworks.

Refer to the following topics when planning your deployment:

- General Requirements
- Install the VM-Series Plugin on Panorama
- Install the Panorama Plugin for GCP
- Prepare to Deploy from the GCP Marketplace

General Requirements

The components in this checklist are common to deploying a VM-Series firewall that you manage directly or with Panorama. Additional requirements apply for Panorama plugin for services such as Stackdriver monitoring, VM monitoring, auto scaling or securing Kubernetes deployments.

Always consult the Compatibility Matrix for Panorama plugin information for public clouds.This release requires the following software:

- **GCP account**—You must have a GCP user account with a linked email address and you must know the username and password for that email address.
- **Google Cloud SDK**—If you have not done so, install Google Cloud SDK, which includes Google Cloud APIs, gcloud and other command line tools. You can use the command line interface to deploy the firewall template and other templates.
- **PAN-OS on VM-Series firewalls on GCP**—VM-Series firewalls running a PAN-OS version available from the Google Marketplace.
  - **VM-Series firewalls**—VM-Series firewalls that you want to manage from Panorama must be deployed in Google Cloud Platform using a Palo Alto Networks image from the Google Marketplace. Firewalls must meet the Minimum System Requirements for the VM-Series Firewall.
  - **VM-Series Licenses**—You must license a VM-Series firewall to obtain a serial number. A serial number is required to add a VM-Series firewall as a Panorama managed device. If you are using the Panorama plugin for GCP to deploy VM-Series firewalls you must supply a BYOL auth code. The Google Marketplace handles your service billing, but the firewalls you deploy will directly interface with the Palo Alto Networks licensing server.
  - **VM-Series plugin on the firewall**—VM-Series firewalls running PAN-OS 9.0 and later include the VM-Series plugin, which manages integration with public and private clouds. As shown

 ©2024 Palo Alto Networks, Inc.

in the Compatibility Matrix, the VM-Series plugin has a minimum version that corresponds to each PAN-OS release.

When there is a major PAN-OS upgrade the VM-Series plugin version is automatically upgraded. For minor releases it is up to you to determine whether a VM-Series plugin upgrade is necessary, and if so, perform a manual upgrade. See Install the VM-Series Plugin on Panorama.

- **Panorama running in Management mode**—A Panorama physical or virtual appliance running a PAN-OS version that is the same or later than the managed firewalls. Virtual instances do not need to be deployed in GCP.
  - You must have a licensed version of Panorama.
  - Panorama must have network access to the VPCs in which the VMs you want to manage are deployed.
  - If you intend to manage VMs deployed in GCP, or configure features such as auto scaling, your PAN-OS and VM-Series plugin versions must meet the Public Cloud requirements to support the Panorama plugin for GCP.
  - VM-Series plugin on Panorama. See Install the VM-Series Plugin on Panorama
- **Panorama plugin for GCP version 2.0.0**—The GCP plugin manages the interactions required to license, bootstrap and configure firewalls deployed with the VM Monitoring or Auto Scaling templates. The GCP plugin, in conjunction with the VM Monitoring or Auto Scaling templates, uses Panorama templates template stacks, and device groups to program NAT rules that direct traffic to managed VM-Series firewalls.

  See Install the Panorama Plugin for GCP.

Install the VM-Series Plugin on Panorama

On Panorama, install or upgrade to the VM-Series plugin version that supports the GCP features you want to configure, as detailed in the Compatibility Matrix table for Public Clouds.

Initial installation—Because the VM-Series plugin is optional on Panorama, the first time you install you must download the VM-Series plugin from the Support portal, then go to **Panorama** > **Device Deployment** > **Plugins** to upload and install.

Upgrade—Go to **Panorama** > **Device Deployment** > **Plugins** and click **Check Now**. Install a version that meets the requirements in the Compatibility Matrix table for Public Clouds.

Install the Panorama Plugin for GCP

The Panorama plugin for GCP is required if you want to use Panorama to manage VM Monitoring or Auto Scaling deployments created with Palo Alto Networks templates. Install the plugin version that supports the GCP features you want to configure, as detailed in the Compatibility Matrix table for Public Clouds.

You cannot upgrade the Panorama Plugin for GCP from version 1.0.0 to version 2.0.x. If you have installed version 1.0.0, remove it before installing 2.0.x.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information

 ©2024 Palo Alto Networks, Inc.

to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Verify your Panorama installation.

On Panorama, ensure that your PAN-OS version meets the requirements to support GCP auto scaling.

STEP 2 | Remove the Panorama plugin for GCP v1.0.

If you have the Panorama plugin v1.0 installed you must remove it.

STEP 3 | Install the Panorama plugin for GCP.

Select **Panorama** > **Plugins**, and type **`gcp`** in the search bar. **Install** the plugin version that supports the features you want to configure (see the Compatibility Matrix table for Public Clouds).

After the installation you can see the plugin in the Panorama dashboard **General Information** list. View **Panorama** > **Google Cloud Platform** and you see the **Setup**, **Monitoring Definition**, and **AutoScaling** interfaces.

STEP 4 | (Optional) If your Panorama appliances are in a high availability configuration, you must manually install the same version of the Google plugin on both Panorama peers.

Configure the Google plugin on the active Panorama peer only. On commit, the configuration syncs to the passive Panorama peer. Only the active Panorama peer polls Google VMs you have configured for VM Monitoring.

Prepare to Deploy from the GCP Marketplace

Review these requirements to ensure that you have proper accounts and permissions before you use the Google Marketplace to deploy the firewall on a Google Compute Engine (GCE) instance.

- General Accounts and Permissions
- Available Google Resources

 ©2024 Palo Alto Networks, Inc.

- Google Authentication Methods
- SSH Key Pair

General Accounts and Permissions

- You, and any users you allow, must have the following minimal roles or equivalent Identity and Access Management (IAM) permissions to connect to the VM-Series firewall:
  - **Compute Viewer**—Compute Viewer enables you to get and list compute engine resources without being able to read the data stored on those resources.
  - **Storage Object Viewer**—Enables you to bootstrap using a Google storage bucket in the same project.

  Users in your organization might have IAM permissions or predefined roles that are more permissive than required. Ensure that you appropriately restrict VM-Series firewall access.

  You can also restrict access with service accounts, as described in Google Authentication Methods.
- **Monitoring Metric Writer**—Required for Stackdriver.

Available Google Resources

Your project must have sufficient resources to deploy the VM-Series firewall as a Google Compute Engine instance. If you are deploying a GCP Marketplace solution, determine whether the solution deploys other VMs in addition to the firewall. In the Google Cloud Console, select **IAM & admin** > **Quotas** to review the resource quotas for your project and the networks and disk space consumed. If you are running out of resources you can ask Google to allocate more for your organization.

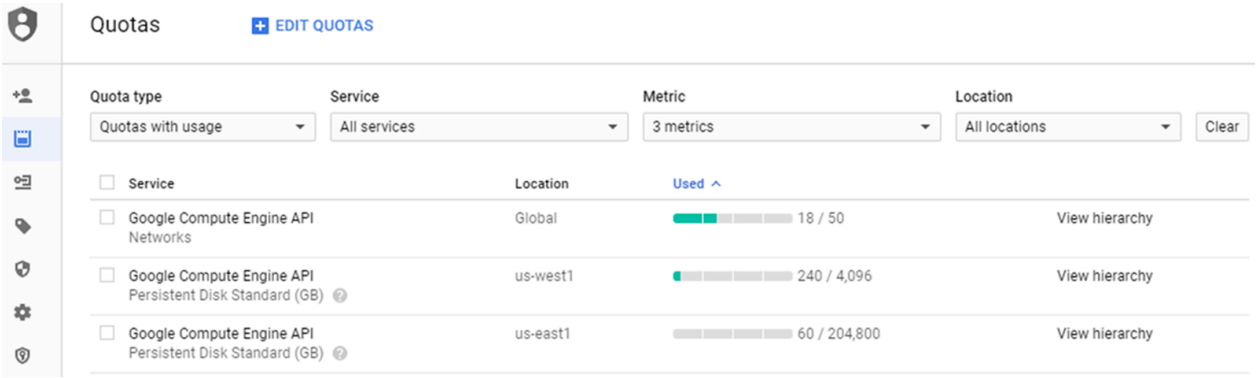

Google Authentication Methods

GCP supports multiple ways to connect to an instance. You can authenticate with a service account or an SSH key pair.

1. **Service Accounts**—Service Accounts apply to applications or VMs—not to end users. They are commonly used to control access when you use programs or scripts, or when you access the firewall from the gcloud command line. If you are using Google Service Accounts to

 ©2024 Palo Alto Networks, Inc.

authenticate instances or applications, you must know the email address for the account(s). Refer to Creating and Managing Service Account Keys.

Using a service account is necessary if you want to connect to the VM-Series firewall from outside the project—either from a different project or from the command line. For example, if you want to enable a physical next generation firewall to monitor your VM-Series firewall, you must save the VM-Series firewall service account information to a JSON file. In the physical firewall, you upload the file when you configure the connection.

1. Select **IAM & Admin** > **Service accounts** and choose **+Create Service Account**.

   Enter the service account name and description, and click **Create**.
2. Select a role type from the drop menu, and on the right, select an appropriate access level.

   For example, select Project > Editor. You can select multiple roles for a service account. When you are finished, click **Continue**.
3. Grant specific users permission to access this service account. Select members from the **Permissions** column on the right to give them permission to access the roles in the previous step.

2. **SSH Keys**—If you deploy the VM-Series firewall from the Marketplace, you must supply one Open SSH key in RSA format for the Google Compute Engine instance metadata.

   The VM-Series firewall only accepts one key at deployment.

   At deployment time, you paste the public key into the Marketplace deployment, as described in SSH Key Pair. After deployment you use the private key to SSH in to the firewall to configure the administrator account. To add users, see Manage Firewall Administrators.

You can authenticate in several ways:

- **Create service accounts for instances**—You can create a service account for a specific instance or instance group, and grant specific permissions, which in turn can be granted to users.
- **Use the default service account for your project**—If you are using the Google Cloud Platform (GCP™) Console, then you logged in with your email address and can access a GCE instance based on whatever permissions or roles the project administrator assigned to your account.

  Every Google Compute Engine instance created with the Google Cloud Console or the gcloud command line tool has a default service account with the name in email address format:

  `<project-number>-compute@developer.gserviceaccount.com`

  To see the service account name for the firewall instance, view the instance details and scroll to the bottom (refer to the Compute Engine default service account).

  The default service account can manage authentication to VMs in the same project as a VM-Series firewall. Access scopes allow the firewall to initiate API calls to VMs in the Google Cloud project.

 ©2024 Palo Alto Networks, Inc.

- **Use IAM permissions and the Google APIs**—If you use the Google SDK APIs and gcloud, then you must call the APIs to authenticate.
  - You typically use the Google SDK when you want to manage the firewall from a command line or you want to run a script to configure the firewall.
  - You need to access the Google APIs if a virtual machine you connect to has a custom image with applications that require Google APIs.

SSH Key Pair

When you deploy the VM-Series firewall from the Google Marketplace you need an SSH key pair to authenticate with the VM-Series firewall.

Create the key pair according to your key generator documentation. Do not edit the public key file. Editing risks introducing illegal characters.

The VM-Series firewall manages authentication differently than GCE instances. After deployment, you first log in with the **admin** user. The VM-series firewall default user name is accepted only once. After a successful login you set an administrator username and password for the VM-Series web interface (see Deploy the VM-Series Firewall from Google Cloud Platform Marketplace).

The Google Marketplace deployment interface **SSH key** field displays the following placeholder:

admin:ssh-rsa your-SSH-key

admin is the VM-Series firewall Administrator user name required to log in to the firewall for the first time. You add the **admin:** prefix into the Marketplace field when you Deploy the VM-Series Firewall from Google Cloud Platform Marketplace.

You cannot log in to the VM-Series firewall if you do not supply the entire public key, or your key has illegal characters when you paste the key into the Marketplace **SSH key** field. When you SSH in to the VM-Series firewall for the first time, the public key is transferred to the firewall.

If the public key is corrupted, you must delete the deployment and start over. Any networks and subnetworks remain, but the firewall rules must be recreated.

STEP 1 | Create an SSH key pair and store the SSH Key pair in the default location for your operating system mentioned in Locating an SSH key.

- **Linux or MacOS**—Use `ssh-keygen` to create the key pair in your .ssh directory.
- **Windows**—Use PuTTYgen to create the key pair.

  The content of the **Key comment** field does not matter to the VM-Series firewall; you can accept the default (the key creation date) or enter a comment that helps you remember the name of the key pair. Use the **Save private key** button to store the private key in your .ssh directory.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Select the full public key.

- **Linux or MacOS**—Open your public key in a text editor and copy the public key.
- **Windows**—You must use the PuTTY Key Generator to view the public key. Launch PuTTYgen, click **Load**, and browse to private key you saved in your .ssh directory.

  In PuTTYgen, scroll down to ensure you select the entire key, right click, and choose Copy.

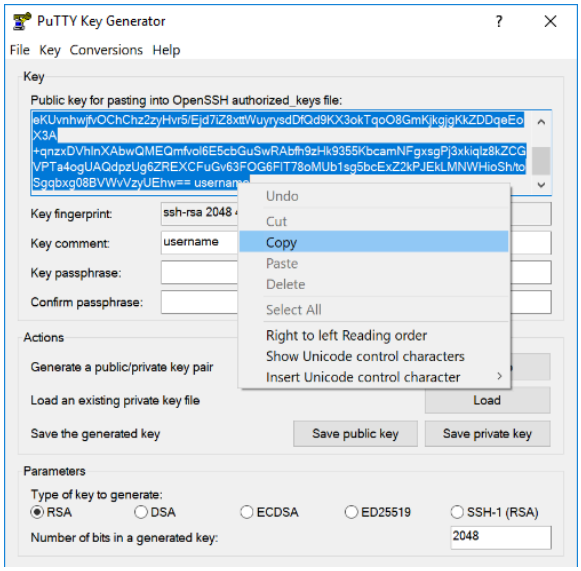

STEP 3 | Enter the public key in the SSH key field as detailed below.

1. In the Marketplace **SSH key** field, delete the placeholder text, and type:

   `admin:`

   Make sure there are no extra spaces following the colon.
2. Insert the cursor after **`admin:`** and choose **Paste as plain text**. The key must be on a single line, as shown below:

   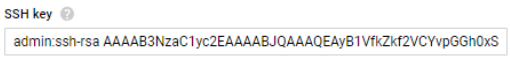

3. Move the cursor to the end of the key, add a space, and type: **`admin`**

   The final contents of the **SSH key** field must be:

   `admin:ssh-rsa` [KEY] **`admin`**

STEP 4 | Check the key.

After the deployment, and before you attempt to log in to the firewall, view the management instance and check the key for linefeeds or extra spaces:

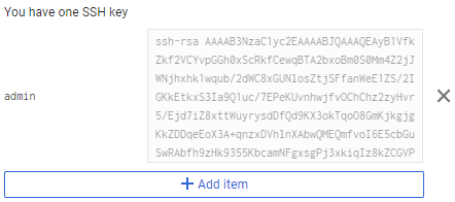

If the key is all on one line and the format is **`admin:ssh-rsa`** [KEY] **`admin`**, you are finished.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | (optional) If something is wrong you must replace the key.

1. Click the **X** to delete the key, then click **+ Add item**.
2. Enter the key as described in Step 3. Now the **SSH key** field must show:

   `admin:ssh-rsa` [KEY] **`admin`**
3. Click **Save** to deploy the updated deployment.
4. Re-check the key.

Virtual Private Cloud (VPC) Network Planning

Before you deploy from the Google Market place, make a plan for VPC networks (referred to as *networks*), subnetworks (also called *subnets*), and Google firewall rules. You must create networks and subnetworks before you start to Deploy the VM-Series Firewall from Google Cloud Platform Marketplace.

The Marketplace deployment page displays only networks and subnetworks that exist when you start the deployment. If a network is missing, you must exit the deployment, create the network, and start over.

- **VPC networks**—You must create a custom network specifically for each VM-Series firewall network interface.
  - See VM-Series Firewall Licenses for Public Clouds to determine the number of network interfaces needed based on your VM-Series firewall license. At a minimum, set up the three VPC networks and subnets required to launch the VM-Series firewall.
  - A GCP project has a default network with preset configurations and firewall rules; you can delete the default network, if unused.
  - By default, there are up to five networks in a project. Your GCP administrator can request additional networks for your project.
  - To connect to the management interface you must create a GCP firewall rules that allows access. You can do this during the deployment if you choose **Enable GCP Firewall rule for connections to Management interface** then supply a CIDR block for **Source IP in GCP Firewall rule for connections to Management Interface**.

    Be sure your networks include all instances you want to secure.
- **Subnetworks**—A compute engine instance can support up to eight Layer 3 interfaces on a single instance. The Management, Trust, and Untrust interfaces consume three interfaces and you can create up to five additional dataplane interfaces. Typically the dataplane interfaces represent application networks.
- **IP address**—You supply IP address ranges when you create interface subnetworks, and you have the option to enable an external address when you deploy a subnetwork.
  - When you create a network subnet, you must specify an IP address range. This range is used for your internal network, so it cannot overlap with other subnets.
  - During deployment, you can choose to enable an external IP address when you create a network interface. By default, you are given an ephemeral IP address. You cannot supply a reserved static IP address during the deployment, but you can promote the

 ©2024 Palo Alto Networks, Inc.

ephemeral address to a static IP address after you complete the deployment process (see Promoting an ephemeral external IP address).

Network Interface Planning

When you deploy from Google Cloud Platform Marketplace, the default VM-Series firewall deployment has three interfaces: the Management plane interface and the Untrust and Trust dataplane interfaces. You can define additional dataplane instances, depending on the available compute resources on your VM; see VM-Series Firewall Licenses for Public Clouds.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

During the deployment you have the opportunity to name these interfaces.

Interface Order

When you deploy with Marketplace, the order of the network interfaces is predefined. The Management interface maps to eth0, Untrust to eth1, and Trust to eth2. Marketplace uses this order because mapping the Management interface to eth0 and the Untrusted interface to eth1 is a requirement if you need to Swap the Management Interface for load balancing.

Management Interface

The first network interface you add is mapped to eth0 on the firewall and includes the option to enable IP forwarding. You use this network interface to manage the VM-Series firewall. Typically, this interface has an external IP address.

An external IP address is only required if a dataplane interface is attached to the public subnet. At creation time, you can receive an ephemeral IP address and later promote it to a static IP address after you complete the deployment (refer to Promoting an ephemeral external IP address).

Dataplane Interfaces (Untrust, Trust)

When you deploy from Marketplace, the order in which you add interfaces is predetermined.

- You configure the Untrust interface after the Management interface. This order means that the untrusted interface is mapped to eth1. The Untrust interfaces are typically attached to the public subnet, and have an external IP address.

  An external IP address is only required if a dataplane interface is attached to the public subnet. At creation time, you can receive an ephemeral IP address, then promote it to a static IP address, as discussed in Promoting an ephemeral external IP address.

- The Trust interface follows the Untrust interface, and it is mapped to eth2. The Trust network often does not have an external IP address. You can add any additional dataplane interfaces after the Trust interface.

Additional Dataplane Interfaces

Plan interfaces for applications you must secure, such as web servers, databases, and other applications in your network. You can create up to five additional dataplane interfaces in addition

 ©2024 Palo Alto Networks, Inc.

to the three required to launch your firewall. Ensure that the applications you want to secure are in networks that connect to the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

Create a Custom VM-Series Firewall Image for Google Cloud Platform

Palo Alto Networks posts VM-Series firewall base image versions or minor versions with critical fixes (such as PAN-OS 11.0) on the Google Cloud Platform (GCP) Marketplace. These versions are available when you deploy a VM-Series firewall from the GCP Marketplace. However, you might need to deploy a PAN-OS version that is earlier or later than the Marketplace version.

To deploy a VM-Series firewall version that is not available on the Marketplace, you can create a custom VM-Series firewall image with a BYOL license.

The basic steps to create a custom firewall from a firewall instance are as follows:

- Deploy a new firewall from the GCP Marketplace.
- Activate your firewall license, download your desired PAN-OS software version to your firewall, use Dynamic Update to update your **Applications and Threats** content, and deactivate the firewall license.
- Perform a private data reset from the GCP console.
- Create a custom image from the upgraded firewall.

STEP 1 | Before you create your custom image, review your accounts, plan and create the networks for VM-Series firewall deployment, for the VM-Series firewall deployment, and plan your network interfaces.

STEP 2 | Deploy the VM-Series firewall from the GCP Marketplace.

You cannot create an image from an existing firewall. Starting from the GCP Marketplace ensures that your custom image can be licensed.

STEP 3 | (BYOL Only) Activate the license.

1. Select **Device** > **Licenses** and activate the licensel.

   The firewall reboots when licensing is complete.
2. Log in to the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Upgrade to your preferred PAN-OS version and install software updates.

1. Select **Device** > **Software** > **Check Now** and download your required PAN-OS version.

   If you do not see the version you want, download it from the Palo Alto Networks customer support website as follows.

   1. Log in and select **Updates** > **Software Updates**.

      From the **Filter By** list, choose PAN-OS for VM-Series.
   2. Select a PAN-OS version and download it to your local machine.
   3. On your VM-Series firewall, **Select Device** > **Software** and **Upload** your PAN-OS version from your local machine to your device.
2. Install your chosen version.
3. Upgrade the PAN-OS software version.
4. Select **Device** > **Dynamic Updates** and upgrade your **Applications and Threats** and any other content you want to include in your base image.

STEP 5 | (BYOL Only) Deactivate the VM from the firewall.

If you do not deactivate the license, you lose the license that you applied on your firewall instance.

1. Select **Device** > **Licenses** and under **License Management**, select **Deactivate VM**.
2. Select **Complete Manually**, and **Export** the license token.
3. Return to the Palo Alto Networks customer support website, select **Assets** > **VM-Series Auth-Codes** > **Deactivate License(s)** and upload the license token.

STEP 6 | Perform a private data reset.

A private data reset removes all logs and restores the default configuration.

The system disks are not erased, so the content updates from Step 4 are intact.

1. Access the firewall CLI and keep it active.
2. From the GCP console, delete SSH keys from your VM-Series firewall.

   1. Select **Compute Engine** > **VM Instances** and select your instance name.
   2. In the **Details** view, select **EDIT**.
   3. Under **SSH Keys**, click the **Show and edit** link and click **X** to remove any SSH keys.
   4. **Save** your changes.
3. (Optional) Export a copy of the configuration.
4. In the CLI, request a private data reset.

   ```
   request system private-data-reset
   ```

   Enter **y** to confirm.

   The firewall reboots to initialize the default configuration.
5. From the GCP console, select **Compute Engine** > **VM instances** and **STOP** the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Create a custom image in the GCP console.

These steps are based on Creating, deleting, and deprecating custom images.

1. Select **Compute Engine** > **Images** > **Create Image**.
2. Name your image and select the **Google-managed key** (see Google-managed encryption keys).
3. Select **Disk** for the Source, and for the **Source disk**, select your stopped VM-Series firewall VM and click **Create**.
4. (Optional) When the image is complete, click the **Equivalent REST** link, and from the **REST response**, copy the selfLink. This is the URI link for any type of CI/CD pipeline that you require.

   For example: `projects/my-vpc-vpcID/global/images/pa-vm-8-1-9`

   Using this link points directly to your image so you can use it in a template or a script. For example:

```
sourceImage: https://www.googleapis.com/compute/v1/projects/
{{project}}/global/images/pa-vm-8-1-9}
```

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on Google Cloud Platform

To deploy the VM-Series firewall using the GCP market place template, you must first create a VPC network for each interface on the firewall. After you deploy the firewall from the Google Marketplace, you can log in to the firewall to adjust the configuration to work within your GCP VPC configuration. You can also enable monitoring so you can collect metrics that enable you to improve resource management or create Security policy rules that automatically adapt to changes in your application environment.

- Deploy the VM-Series Firewall from Google Cloud Platform Marketplace
- Management Interface Swap for Google Cloud Platform Load Balancing
- Use the VM-Series Firewall CLI to Swap the Management Interface
- Enable Google Stackdriver Monitoring on the VM Series Firewall
- Enable VM Monitoring to Track VM Changes on Google Cloud Platform (GCP)
- Use Dynamic Address Groups to Secure Instances Within the VPC
- Use Custom Templates or the gcloud CLI to Deploy the VM-Series Firewall

Deploy the VM-Series Firewall from Google Cloud Platform Marketplace

You can use Google® Cloud Platform Marketplace to deploy the VM-Series firewall on a fixed vCPU capacity license (VM-Series Models). The licensed images available from public clouds are:

- VM-Series Next-Generation Firewall Bundle 1
- VM-Series Next-Generation Firewall Bundle 2
- VM-Series Next-Generation Firewall (BYOL)

See Deploy the VM-Series Firewall from Google Cloud Platform Marketplace for more about these license options.

The Marketplace deploys an instance of the VM-Series firewall with a minimum of one management interface and two dataplane interfaces (Trust and Untrust). You can add additional dataplane interfaces for up to five Google Compute Engine instances in your virtual private cloud (VPC).

Before you deploy the VM-Series firewall, you must create or choose a project in your organization and create any networks and subnets that will connect to the firewall, as described in VPC Network Planning and Network Interface Planning.

You cannot attach multiple network interfaces to the same VPC network. Every interface you create must have a dedicated network with at least one subnet. Ensure that your networks include any additional dataplane instances you create.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Choose a Bootstrap Method.

STEP 2 | Locate the VM-Series firewall listing in the Marketplace.

1. Log in to the Google Cloud Console.
2. From the Products and Services menu, select **Marketplace**.
3. Search for **`VM-Series`**.
4. Select one of the VM-Series firewall licensing options.

STEP 3 | Click **Launch on Compute Engine**.

STEP 4 | Name the instance and choose resources.

1. Enter the **Deployment Name** (this name is displayed in the Deployment Manager). The name must be unique and cannot conflict with any other deployment in the project.
2. Select a **Zone**. See Regions and Zones for a list of supported zones.
3. Select a **Machine Type** based on the VM-Series System Requirements for your license and the Minimum System Requirements for the VM-Series Firewall on Google Cloud Platform.

STEP 5 | **Specify instance metadata**.

The options **Bootstrap Bucket** and **Interface Swap** affect the initial configuration the first time the VM-Series firewall boots.

1. **Bootstrap Bucket** (Optional)—If you plan to use a bootstrap file, enter the name of a storage bucket, or the path to a folder within the storage bucket, that contains the bootstrap package. You need permission to access the storage bucket. For example:

   ```
   vmseries-bootstrap-gce-storagebucket=<bucketname>
   ```

   or

   ```
   vmseries-bootstrap-gce-storagebucket=<bucketname/
   directoryname>
   ```

   If you choose to bootstrap with custom metadata, continue to Step 6.
2. **Interface Swap** (Optional)—Swap the Management interface (eth0) and the first dataplane interface (eth1) at deployment time. Interface swap is only necessary when you deploy the VM-Series firewall behind Google Cloud Platform HTTP(S) Load

 ©2024 Palo Alto Networks, Inc.

Balancing. For details, see Management Interface Swap for Google Cloud Platform Load Balancing.

3. **SSH key**—Paste in the public key from an SSH key pair. Follow the instructions for your OS in SSH Key Pair, to create, copy, and paste the key. Windows users must view the key in PuTTY, copy from the user interface, and paste into Marketplace deployment.

   If the key is not formatted properly, the VM-Series firewall does not allow you to log in. You must delete the deployment and start over.

4. Click **More** to reveal additional metadata options. The options **blockProjectKeys**, and **enableSerialConsole** are properties of the instance; you can change these metadata values after a successful deployment.

   - **blockProjectKeys** (Optional)—If you Block Project Keys, you can use only the public SSH key you supply to access the instance.
   - **enableSerialConsole** (Optional)—Interacting with the Serial Console enables you to monitor instance creation and perform interactive debugging tasks.

STEP 6 | **Specify custom metadata**.

If you choose to bootstrap with custom metadata, add any key-value pairs that you did not add in Step 5. See init-cfg.txt File Components for the list of key-value pairs. For example:

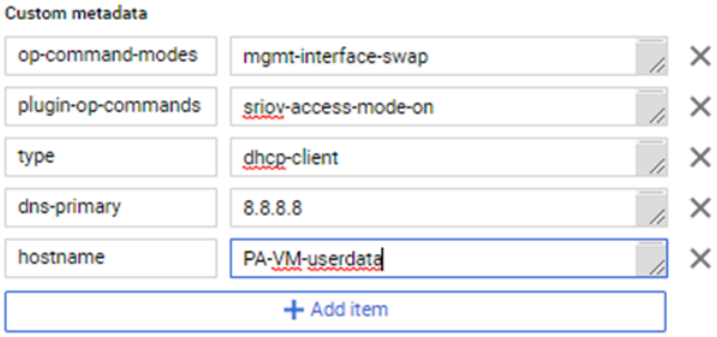

STEP 7 | **Configure the boot disk**.

1. **Boot disk type**—Select from SSD Persistent disk or Standard Persistent Disk. See Storage Options.
2. Enter the **Boot disk size**—60GB is the minimum size. You can edit the disk size later but you must stop the VM to do so.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Configure the management interface.

1. **Management VPC Network name**—Choose an existing network
2. **Management Subnet name**—Choose an existing subnet.
3. **Enable External IP for Management interface** (Optional)—If you enable this option, you can use the IP address assigned to the VM-Series firewall management interface to use SSH to access the VM-Series firewall web interface.
4. **Enable GCP Firewall rule for connections to Management interface** (Optional)—This option automatically creates a GCP firewall Allow rule for an external source IP address that you supply.
5. **Source IP in GCP Firewall rule for connections to Management Interface**—If you **Enable GCP Firewall rule for connections to Management interface**, enter a source IP address or a CIDR block.
   - Do not use 0.0.0.0/0. Supply an IP address or a CIDR block that corresponds to your dedicated management IP addresses or network. Do not make the source network range larger than necessary.
   - Verify the address to ensure that you do not lock yourself out.

STEP 9 | Configure the Untrust dataplane interface.

1. **Untrust VPC Network name**—Choose an existing network.
2. **Untrust Subnet name**—Choose an existing subnet.
3. **Enable External IP for Untrust**—Enable GCP to provide an ephemeral IP address to act as the external IP address.

STEP 10 | Configure the Trust dataplane interface.

1. **Trust VPC Network name**—Choose an existing network.
2. **Trust Subnet name**—Choose an existing network.
3. **Enable External IP for Trust**—Enable GCP to provide an ephemeral IP address to act as the external IP address.

STEP 11 | Configure additional interfaces. You must enter the number of dataplane interfaces you want to add; the default is 0 (none). The deployment page always displays fields for five additional dataplanes numbered 4 through 8.

1. **Additional Dataplane interfaces**—Enter the number of additional dataplane instances.

   If this number is 0 (default), dataplane numbers 4 through 8 are ignored even if you fill out the interface fields. If, for example, you specify 2 and then fill out information for three interfaces, only the first two are created.

2. **Additional Dataplane # VPC name**—Choose an existing network.
3. **Dataplane # Subnet name**—Choose a subnet that exists.
4. **Enable External IP for dataplane # interface**—Enable GCP to provide an ephemeral IP address to act as the external IP address.

STEP 12 | **Deploy** the instance.

STEP 13 | Use Google Cloud Deployment Manager to view and manage your deployment.

 ©2024 Palo Alto Networks, Inc.

STEP 14 | Use the CLI to change the administrator password on the firewall.

1. Log in to the VM-Series firewall from the command line. In your SSH tool, connect to the External IP for the management interface, and specify the path to your private key.

   Windows users: Use PuTTY to connect to the VM-Series firewall and issue command line instructions. To specify the path to the private key, select **Connection** > **SSH** > **Auth**. In **Private key file for authentication**: click **Browse** to select your private key.
2. Enter configuration mode:

   ```
   VMfirewall> configure
   ```
3. Enter the following command:

   ```
   VMfirewall# set mgt-config users admin password
   ```
4. Enter and confirm a new password for the administrator.
5. Commit your new password:

   ```
   VMfirewall# commit
   ```
6. Return to command mode:

   ```
   VMfirewall# exit
   ```
7. (Optional) If you used a bootstrap file for interface swap, use the following command to view the interface mapping:

   ```
   VMfirewall> debug show vm-series interfaces all
   ```

STEP 15 | Access the VM-Series firewall web interface.

1. In a browser, create a secure (https) connection to the IP address for the management interface.

   If you get a network error, check to see that you have a GCP firewall rule that allows the connection.
2. When prompted, enter the username (admin) and the administrator password you specified from the CLI.
3. (Optional) If you bootstrapped, then Verify Bootstrap Completion.

   If you see problems, search the log information on the VM-Series firewall. Choose **Monitor** > **System** and, in the manual search field, enter **`description contains 'bootstrap'`** and look for a message in the results that indicates that the bootstrap was successful.

After you log in to the firewall, you can add administrators and create interfaces, zones, NAT rules, and policy rules, just as you would on a physical firewall.

Management Interface Swap for Google Cloud Platform Load Balancing

Because internal load balancing can send traffic only to the primary interface of the next hop load-balanced Google Compute Engine instance, the VM-Series firewall must be able to use eth0 for dataplane traffic.

 ©2024 Palo Alto Networks, Inc.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

The firewall can receive dataplane traffic on eth0 if the VM-Series firewall is behind the Google Cloud Platform internal load balancing interface.

- The VM-Series firewalls secure traffic outbound directly to the internet without requiring a VPN link or a Direct Connect link back to the corporate network.
- The VM-Series firewall secures an internet-facing application when there is exactly one back-end server, such as a web server, for each firewall. The VM-Series firewalls and web servers can scale linearly, in pairs, behind the Google internal load balancing address.

To allow the firewall to send and receive dataplane traffic on eth0 instead of eth1, you must swap the mapping of the internal load balancing network interface within the firewall so that eth0 maps to ethernet 1/1, and eth1 maps to the MGT interface on the firewall.

If possible, swap the management interface mapping before you configure the firewall and define policy rules.

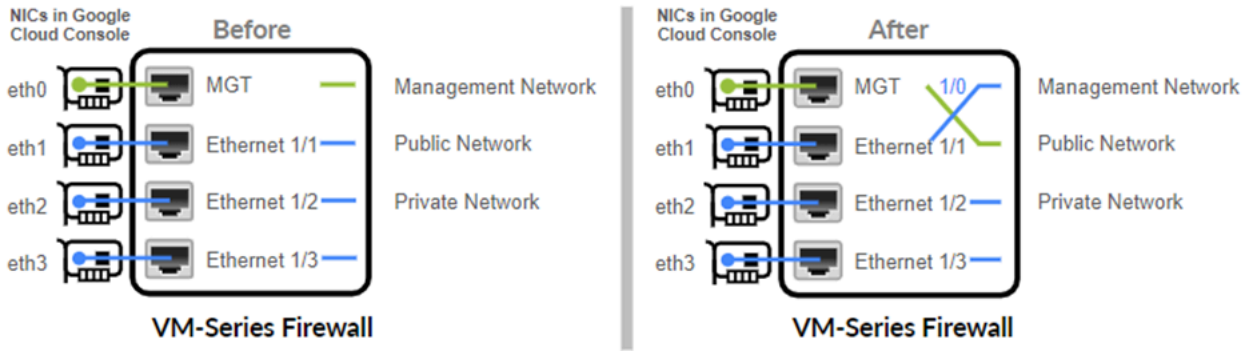

Swapping how the interfaces are mapped allows Google Cloud Platform to distribute and route traffic to healthy instances of the VM-Series firewall located in the same or different zones.

Swap the Management Interface

You can swap the interfaces when you , or you can configure the firewall after it is created.

At creation— When you deploy the VM-Series firewall, you can enable interface swap in two ways.

- Google Cloud Console — In the Create Instance form, enter a key-value pair in the **Metadata** field, where **`mgmt-interface-swap`** is the key, and **`enable`** is the value.
- Bootstrap File — Create a bootstrap file the includes the **`mgmt-interface-swap`** operational command in the bootstrap configuration, as described in Bootstrap the VM-Series Firewall on Google Cloud Platform. In the Create Instance form, enter a key-value pair in the **Metadata** field to enable the bootstrap option.

From the VM-Series firewall—Log in to the firewall, and Use the VM-Series Firewall CLI to Swap the Management Interface. In operational mode, issue the following command:

```
set system setting mgmt-interface-swap enable yes
```

 ©2024 Palo Alto Networks, Inc.

- *Pick one method to specify the interface swap setting— the bootstrap configuration file, the firewall CLI, or the Google Compute Engine instance **Metadata** field (accessed from the Google Cloud Console). Using one method ensures predictable behavior on the firewall.*

  From the Google Cloud Console you cannot confirm whether you have swapped eth0 and eth1. After swapping, you must remember that load balancing is on eth0 and the firewall management interface is eth1 so that you can properly configure Google Cloud Platform load balancing, and create security policy rules to secure load balancing to one or more VM-Series firewalls.
- *If you configured the VM-Series firewall before swapping, check whether any IP address changes for eth0 and eth1 impact policy rules.*

Use the VM-Series Firewall CLI to Swap the Management Interface

This task is only required if your architecture places the VM-Series firewall behind the Google Cloud Platform internal load balancer.

If you did not specify metadata to swap the management interface (MGT) with the dataplane interface when you deployed the firewall, you can use the CLI to enable the firewall to receive dataplane traffic on the primary interface.

STEP 1 | Deploy the VM-Series Firewall from Google Cloud Platform Marketplace.

Before you proceed, verify that the firewall has a minimum of two network interfaces (eth0 and eth1). If you launch the firewall with only one interface, the interface swap command causes the firewall to boot into maintenance mode.

STEP 2 | On the Google Cloud Console, view the VM instance details to verify the network interface IP addresses of the eth1 interface and verify that any security rules allow connections (HTTPS and SSH) to the new management interface (eth1).

STEP 3 | Log in to the VM-Series firewall CLI and enter the following command:

```
set system setting mgmt-interface-swap enable yes
```

You can view the default mapping from the command line interface. The output is similar to this:

```
> debug show vm-series interfaces all
Interface_name  Base-OS_port
  mgt              eth0
  Ethernet1/1      eth1
  Ethernet1/2      eth2
```

STEP 4 | Confirm that you want to swap the interface (use the eth1 dataplane interface as the management interface).

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Reboot the firewall for the swap to take effect:

```
request restart system
```

STEP 6 | Verify that the interfaces have been swapped:

```
debug show vm-series interfaces all
```

Enable Google Stackdriver Monitoring on the VM Series Firewall

A VM-Series firewall on a Google® Compute Engine instance can publish custom PAN-OS metrics to Google Stackdriver. These metrics allow you to assess performance and usage patterns so that you can manage your firewall resources accordingly.

- Google Stackdriver Permissions
- Enable Google Stackdriver

Google Stackdriver Permissions

Authentication requirements vary based on whether you can use the default service account to authenticate or need to use Google APIs to authenticate.

You can authenticate in two ways:

- **Use the default service account for the VM-Series firewall instance**—If you are using the Google Cloud Platform (GCP™) Console, then you logged in with your email address and can access the instance based on whatever permissions or roles the project administrator assigned to your account.
- **Use IAM permissions and the Google APIs**—If you use the Google SDK APIs and gcloud, then you must call the APIs to authenticate. You typically use the Google SDK when you want to manage the firewall from a command line or you want to run a script to configure the firewall.

Every Google Compute Engine instance created with the Google Cloud Console or the gcloud command line tool has a default service account with the name in email address format:

`<project-number>-compute@developer.gserviceaccount.com`

To see the service account name for the firewall instance, view the instance details and scroll to the bottom (refer to the Compute Engine default service account).

The default service account can manage authentication for monitoring VMs in the same project as a VM-Series firewall.

- Access scopes allow the firewall to initiate API calls to monitor VMs in a Google Cloud project.
- You don't need to access the Google APIs unless one of the monitored virtual machines has a custom image with applications that require Google APIs.

If you want to set up monitoring from a physical firewall or from a VM-Series firewall in a different project, you must use the Google APIs to authenticate. There are two prerequisites:

- Google APIs must be installed.
- Your account must have the roles Monitoring Metric Writer and Stackdriver Account Viewer.

 ©2024 Palo Alto Networks, Inc.

Enable Google Stackdriver

For a description of the PAN-OS metrics that you can publish to Google Stackdriver, see Custom PAN-OS Metrics Published for Monitoring.

STEP 1 | Push PAN-OS metrics from a VM-Series firewall on a Google Compute Engine instance to Stackdriver.

1. Log in to the web interface on the VM-Series firewall.
2. Select **Device** > **VM-Series**. Under Google Cloud Stackdriver Monitoring Setup, click **Edit** ( ).
   1. Check **Publish PAN-OS metrics to Stackdriver**.

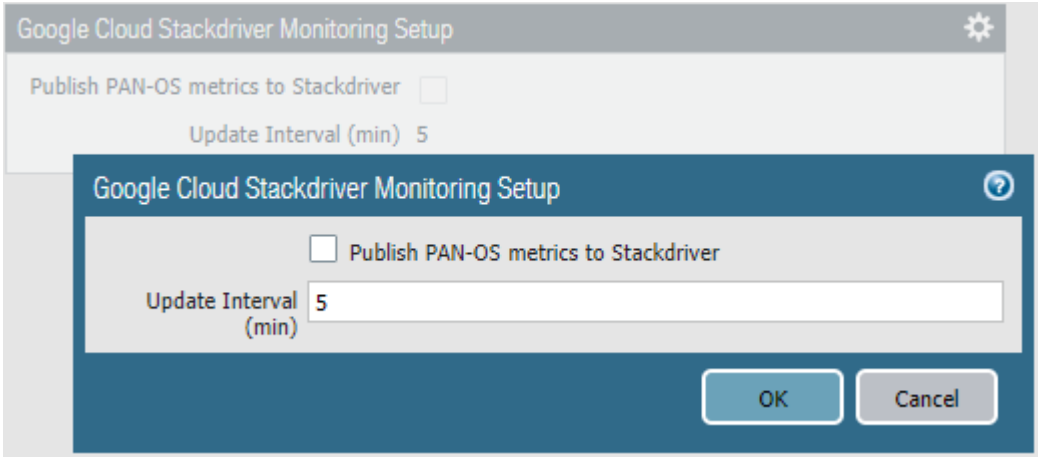

   2. Set the **Update Interval** (range is 1 - 60 minutes; default is 5). This is the frequency at which the firewall publishes the metrics to Stackdriver.
   3. Click **OK**.
3. **Commit** your changes.

   Wait until the firewall starts to publish metrics to Stackdriver before you configure alarms for PAN-OS metrics.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Verify that you can see the metrics on Stackdriver.

1. In the Google Cloud Console, select **Products and Services** > **Monitoring**.
2. In Stackdriver, choose **Resources** > **Metrics Explorer**.
3. In the Find resource type and metric section, enter **`custom`** in the search field to filter the PAN-OS metrics.

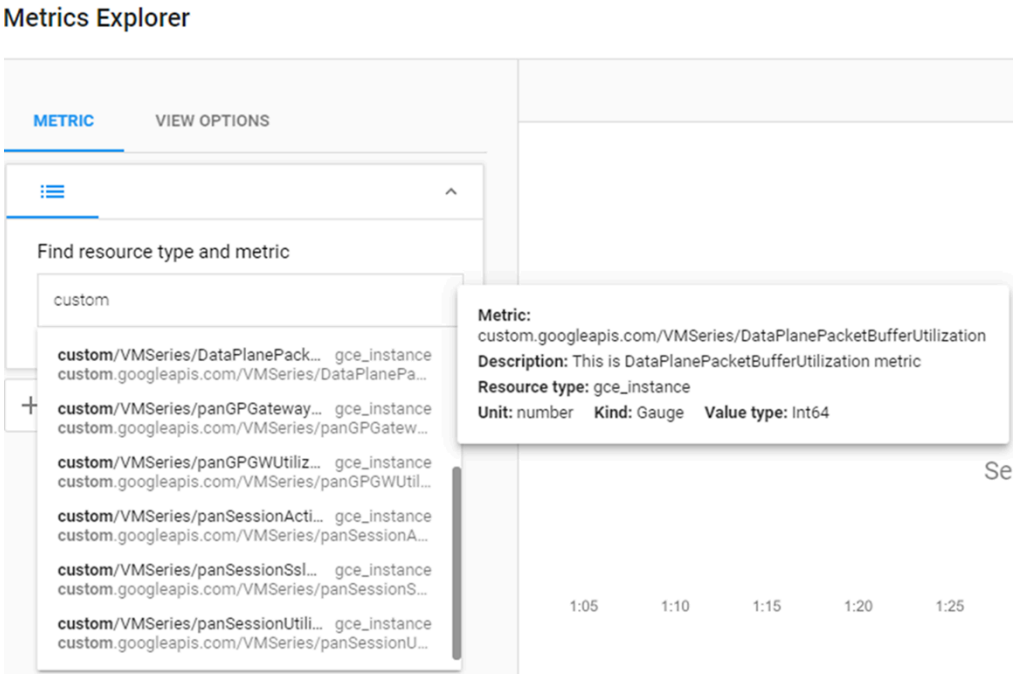

STEP 3 | Configure alerts and actions for PAN-OS metrics on Stackdriver. See Monitoring Quickstart for Google Compute Engine and Stackdriver Introduction to Alerting.

Enable VM Monitoring to Track VM Changes on Google Cloud Platform (GCP)

You can enable any firewall that runs PAN-OS 9.0 (virtual or physical) to monitor application workloads deployed on Google Compute Engine instances. VM Monitoring enables you to monitor a predefined set of metadata elements or attributes on the VM-Series firewall. In the PAN-OS 9.1 Administrator's Guide, see Attributes Monitored on Virtual Machines in Cloud Platforms.

With an awareness of virtual machine adds, moves, and deletes within a Google VPC, you can create Security policy rules that automatically adapt to changes in your application environment. As you deploy or move virtual machines, the firewall collects attributes (or metadata elements). You can use this metadata for policy matching and to define Dynamic Address Groups (see Use Dynamic Address Groups to Secure Instances Within the VPC).

You can configure up to ten VM information sources on each firewall or on each virtual system on a firewall capable of multiple virtual systems. Information sources can also be pushed using Panorama templates.

To perform VM monitoring, you must have the IAM role Monitoring Metric Writer.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Log in to your deployed firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enable VM Monitoring.

1. Select **Device** > **VM Information Sources**.
2. **Add** a VM information source and enter the following information:
   - Specify a **Name** to identify the instance that you want to monitor.
   - Select the Google Compute Engine **Type**.
   - Select **Enabled**.
   - Choose the **Service Authentication Type**.
     - If you choose **VM-Series running in GCE**, you are authenticating with the default service account generated when an instance is created. This is part of the instance metadata.
     - If you want to monitor from a firewall outside the current project, choose **Service Account**. You must upload the service account credentials in JSON format. See Creating and Managing Service Account Keys.
   - (Optional) Modify the **Update interval** to a value between 5-600 seconds. By default the firewall polls every 5 seconds. The API calls are queued and retrieved every 60 seconds, an update takes up to 60 seconds plus the configured polling interval.

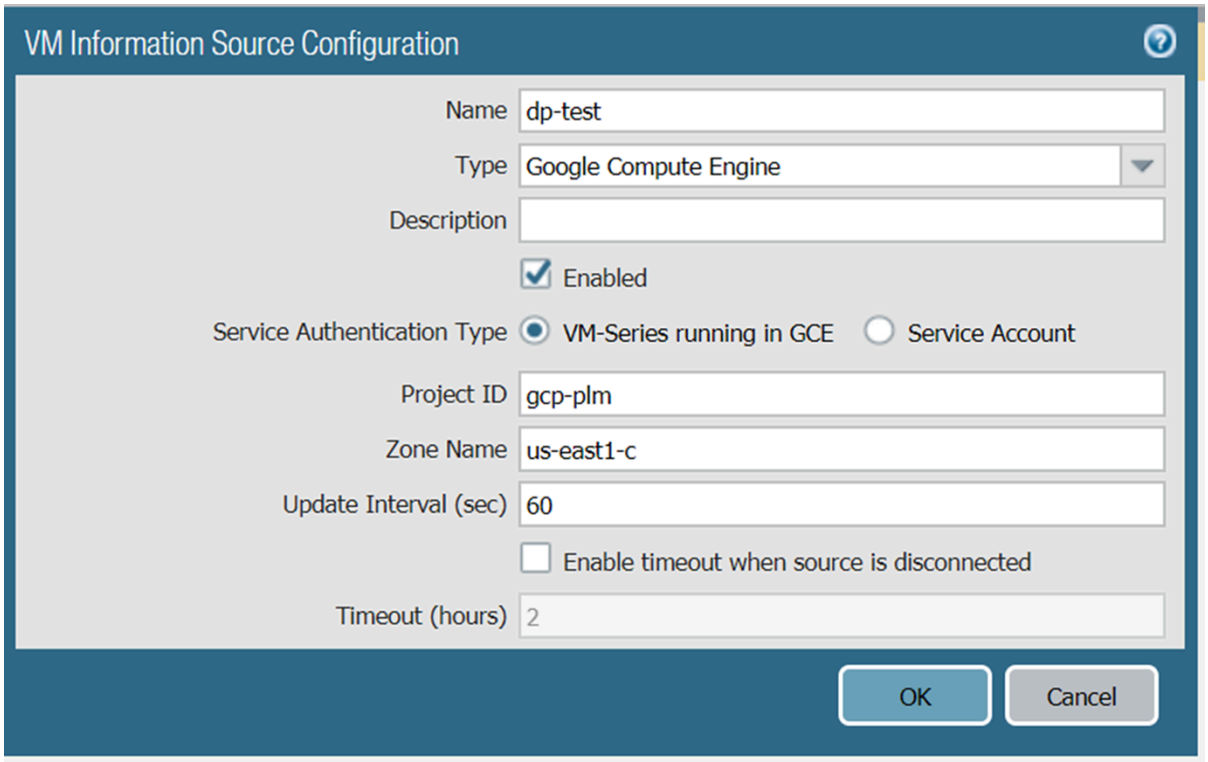

   - (Optional) To change the number of hours before timeout, check **Enable timeout when the source is disconnected** and enter the Timeout (hours) before the connection to the monitored source is closed (range is 2 to 10; default is 2).

     If the firewall cannot access the host and the specified limit is reached, the firewall closes the connection to the source.
   - Click **OK** and **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

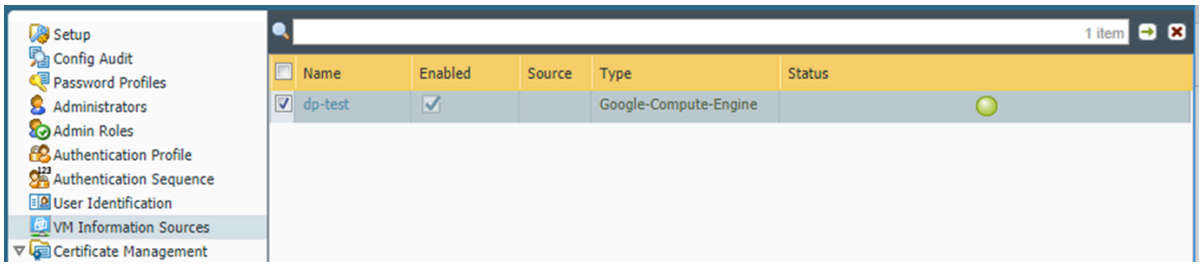

STEP 3 | Verify the connection status.

If the connection status is pending or disconnected, verify that the source is operational and that the firewall is able to access the source. If you use a port other than the Management (MGT) port for communicating with the monitored source, then you must change the service route (select **Device** > **Setup** > **Services**, click **Service Route Configuration**, and modify the **Source Interface** for the **VM Monitor** service).

Use Dynamic Address Groups to Secure Instances Within the VPC

In a dynamic environment such as the Google® Cloud Platform (GCP™), where you launch new instances on demand, the administrative overhead in managing Security policy can be cumbersome. Using use dynamic address groups in policy enables agility and prevents disruption in services or gaps in protection.

This workflow assumes that you have deployed the VM-Series firewall, configured some applications on instances, and enabled Google Stackdriver monitoring.

STEP 1 | Configure the firewall to monitor the VPC.

STEP 2 | Label instances in the VPC.

A label is a name-value pair. You can label resources from the Google Cloud Console, from Google API calls, or from the Google Cloud Shell. In this task we are labeling instances; however, labels can be applied to many resources, as described in Labeling Resources.

You can also add labels from the Instance browser.

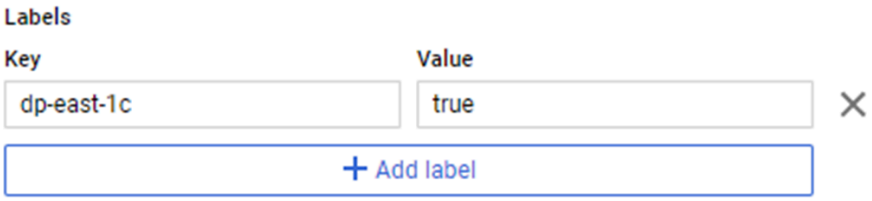

The labels you create support your strategy for differentiating your resources in ways that are useful to your Security policy.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create a dynamic address group on the firewall.

1. Select **Objects** > **Address Groups**.
2. **Add** a dynamic address group and specify a **Name** and a **Description**.
3. Set **Type** to **Dynamic**.
4. Define the match criteria.
   1. **Add Match Criteria** and select the **And** operator.
   2. Select the attributes to filter for or to match against.

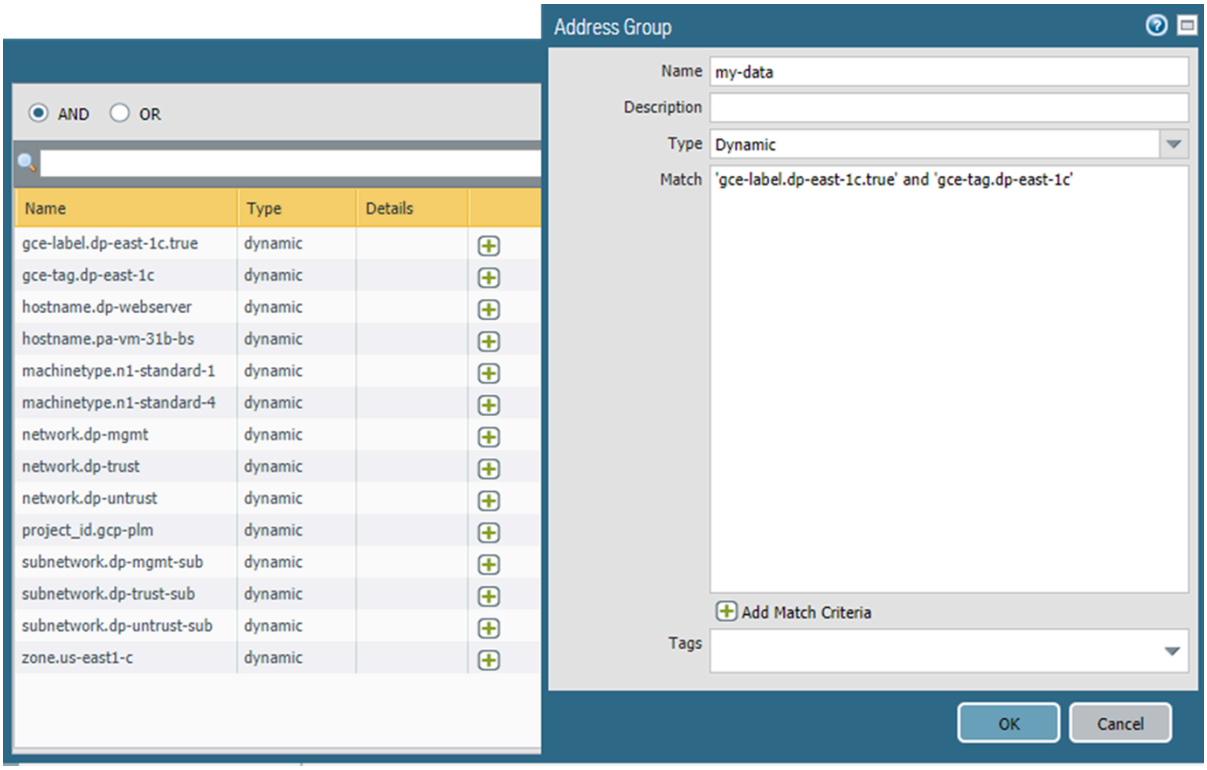

5. Click **OK**.
6. Click **Commit**.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Use the dynamic address group in a Security policy rule.

Create a rule to allow internet access to any web server that belongs to the dynamic address group called my-data.

1. Select **Policies** > **Security**.
2. **Add** a rule and a **Name** for the rule and verify that the **Rule Type** is `universal`.
3. In the **Source** tab, add trust as the **Source Zone**.
4. In the Source Address section, **Add** your new my-data group.
5. In the **Destination** tab, add untrust as the **Destination Zone**.
6. In the **Service/URL Category** tab, verify that the service is set to **application-default**.
7. In the **Actions** tab, set the **Action** to Allow.
8. In the Profile Settings, set the **Profile Type** to **Profiles** and then attach the default profiles for Antivirus, Anti-Spyware, and Vulnerability Protection.
9. Click **OK**.
10. Click **Commit**.

STEP 5 | Verify that members of the dynamic address group are populated on the firewall.

Policy will be enforced for all IP addresses that belong to this address group and that are displayed here.

1. Select **Policies** > **Security** and select the rule.
2. Select **Inspect** from the drop-down. You can also verify that the match criteria is accurate.
3. Click **more** to verify that the list of registered IP addresses is displayed.

Use Custom Templates or the gcloud CLI to Deploy the VM-Series Firewall

The official VM-Series images published on Google Cloud Platform Marketplace are available in the **paloaltonetworksgcp-public** project. You need to know the secure path to these images if you want to call them from the **gcloud** command line, or refer to them in a template you have written or adapted.

- BYOL: vmseries-byol-<*version*>
- PAYG Bundle 1: vmseries-bundle1-<*version*>
- PAYG Bundle 2: vmseries-bundle2-<*version*>

Use the gcloud CLI to find the current image names and project:

```
gcloud compute images list --project paloaltonetworksgcp-public
--no-standard-images
NAME                  PROJECT                     FAMILY  DEPRECATED
STATUS
vmseries-bundle1-810  paloaltonetworksgcp-public
 READY
vmseries-bundle2-810  paloaltonetworksgcp-public
 READY
```

 ©2024 Palo Alto Networks, Inc.

```
  vmseries-byol-810     paloaltonetworksgcp-public
   READY
```

Add the **--uri** flag to see the image paths:

```
  gcloud compute images list --project paloaltonetworksgcp-public
  --no-standard-images --uri
```

```
  https://www.googleapis.com/compute/v1/projects/paloaltonetworksgcp-
public
  /global/images/vmseries-bundle1-810
  https://www.googleapis.com/compute/v1/projects/paloaltonetworksgcp-
public
  /global/images/vmseries-bundle2-810
  https://www.googleapis.com/compute/v1/projects/paloaltonetworksgcp-
public
  /global/images/vmseries-byol-810
```

For an example, download the gcp-two-tier template from https://github.com/PaloAltoNetworks.

This template separates the image name (which includes the PAN-OS version) from the URL path. In `two-tier-template.py` the *image* variable expects the image name; for example: *vmseries-byol-810*. `vm-series-template.py` uses the values of *COMPUTE_URL_BASE* and *sourceImage* to build the path.

 ©2024 Palo Alto Networks, Inc.

VM Monitoring with the Panorama Plugin for GCP

The Panorama plugin for Google Cloud Platform (GCP) version 2.0.0 enables you to create a VM monitoring configuration that authenticates with a GCP project and monitors VM-Series firewalls and other VMs deployed within it. Once you establish a connection to your project, the plugin can retrieve IP-address-to-tag communication between Panorama and GCP assets. Tags can be predefined attributes, user-defined labels for VMs, and user-defined network tags (see Review and Create Tags).

The Panorama plugin for GCP retrieves the internal and external IP addresses from running VMs, and periodically retrieves IP-to-tag mappings from VMs in connected GCP VPCs.

You can use tags to organize VMs into dynamic address groups, and then reference your tags in Security policy rules that allow or deny traffic to specific VM IP addresses. To consistently enforce Security policy, you can then push rules to your VM-Series firewalls.

- Configure VM Monitoring with the Panorama Plugin for GCP

Configure VM Monitoring with the Panorama Plugin for GCP

This topic describes the steps to prepare your GCP assets for VM monitoring, review the required Panorama elements, and describes how to configure VM Monitoring in the Panorama plugin for Google Cloud Platform (GCP).

- Configure GCP Assets for VM Monitoring
- Review and Create Tags
- Configure VM Monitoring with the Panorama Plugin for GCP
  - Prepare Panorama to Configure VM Monitoring
  - Set Up VM Monitoring

Configure GCP Assets for VM Monitoring

You can monitor VM-Series firewalls you deployed from the GCP marketplace, firewalls you deployed with auto scaling Firewall templates, GCE instances you created from to the GCP console or the gcloud command line, or other virtual machines deployed in GCP. If you deploy PAN-OS VMs from the Marketplace, follow the instructions in Set Up the VM-Series Firewall on Google Cloud Platform.

Review IAM Roles

Ensure that you have the following minimum permissions for VM Monitoring tasks:

- In GCP console, create a service account for your project and grant the permission project owner or editor.

  Service account creation cannot be automated. If you do not have permission to create a service account you can ask an administrator to create it and assign an appropriate role to you.
- View your service account: read-only.
- View PAN-OS VMs deployed from the Google Marketplace: Compute viewer.
- Assign a user-defined tag to an instance: Project owner, editor or Instance Admin.

 ©2024 Palo Alto Networks, Inc.

Create a Service Account

Before you use the GCP plugin on Panorama to configure VM Monitoring, you must use the GCP console to create a service account that grants permissions to access your GCP project, VM-Series firewalls deployed within it, any other VMs that you want Panorama to manage, and related networks and subnetworks. The GCP plugin for Panorama retrieves pre-defined attributes for Google assets, user defined VM labels, and user-defined network tags.

Every project has a default service account that was automatically created when the project was created. If you create a separate service account specifically for VM Monitoring you have greater control of users and their roles. You can configure up to 100 service accounts per project.

STEP 1 | In the Google Cloud Platform console, select the project you want to monitor.

STEP 2 | Select **IAM & Admin** > **Service accounts** and choose **+Create Service Account**.

Enter the service account name and description, and click **Create**.

STEP 3 | Select a role type from the drop menu, and on the right, select an appropriate access level.

For example, select Project > Editor. You can select multiple roles for a service account.

When you are finished, click **Continue**.

STEP 4 | Grant specific users permission to access this service account. Select members from the **Permissions** column on the right to give them permission to access the roles in the previous step.

STEP 5 | (Optional) Click **+CREATE KEY** to create a credential that allows you to authenticate with the Google Cloud CLI to access VM-Series firewalls, networks, and other VMs associated with this service account.

The key is downloaded automatically. Be sure to store it in a secure location. The JSON format for the generated private key is as follows:

```
{
  "type": "service_account",
  "project_id": "gcp-xxx",
  "private_key_id": "252e1e7a2e9c84b5d4dbb6195b1de074594b6499",
  "private_key": "-----BEGIN PRIVATE KEY-----
\nMIIEvgIBADANBgkqhkiG9w0BAQEFAASCBKgwggSkAgEAAoIBAQDAd0i+RMKCtrs0
\n4KHnzTAPrgoBjRgpjyNcvQmdUqHr\n-----END PRIVATE KEY-----\n",
  "client_email": "dlp-vm-monit-svc-acct@gcp-
xxx.iam.gserviceaccount.com",
  "client_id": "108932514695821539229",
  "auth_uri": "https://accounts.google.com/o/oauth2/auth",
  "token_uri": "https://oauth2.googleapis.com/token",
  "auth_provider_x509_cert_url": "https://www.googleapis.com/
oauth2/v1/certs",
  "client_x509_cert_url": "https://www.googleapis.com/
robot/v1/metadata/x509/dlp-vm-monit-svc-acct%40gcp-
xxx.iam.gserviceaccount.com"
}
```

 ©2024 Palo Alto Networks, Inc.

Review and Create Tags

"Tag" is a general term for predefined attributes, user-defined labels, and user-defined network tags.

- Predefined tags (attributes) are automatically created for Google VMs. When you configure VM Monitoring you can choose to monitor all 8 of the predefined attributes, or you can create a customized list of attributes to monitor.
- You can define your own tags for VM labels and network tags.

Tag VMs and networks so that you can identify and group them so that you can structure rules to enforce Security policy. You can tag any VM deployed in your Google project—for example, a VM-Series firewall, a web server, an application server, or a load balancer.

- Tags must be associated with a VM. This also applies to networks and subnetworks.
- If there are multiple IP addresses associated with an instance (for example if you tagged the VM-Series firewall trust and untrust interfaces), Panorama generates multiple sets of tag information.

The total number of tags that the Panorama plugin can retrieve and register depends on the PAN-OS version Panorama is running and the version of the managed VM-Series firewalls.

Google zone, Google region, VPC name, and Subnet name are used to tag network interfaces on VMs with multiple interfaces. specific to network interface.

Predefined Attributes

The Google Cloud Platform plugin for Panorama retrieves the following predefined tags from any managed VM:

- **Project ID**—For example: google.project-id.myProjectId.

  To find your project information in the Google console, select your project, then select **IAM & Admin** > **Settings**.
- **Service account**—Your service account in the form of an email address. For example: google.svc-accnt.sa-name@project-id.iam.gserviceaccount.com.

  To find your Service account, view the VM instance details.
- **VPC name**—The name of the VPC network for a managed VM. For example: google.vpc-name.myvnet.
- **Subnet name**—The name of a subnet you created for a managed VM interface. For example, for the VM-Series firewall untrust interface, the name of the subnet you created for the untrust interface: google.subnet-name-untrust.web.
- **OS SKU**—The operating system you chose when you deployed the managed VM. For example: google.os-sku.centos-7.

  This attribute is not supported if the VM uses a custom image.
- **Google zone**—The zone you selected when you deployed the VM. For example: google.zone.us-east1-c.
- **Google region**—The region containing the zone you selected. For example: google.region.us-east1.

 ©2024 Palo Alto Networks, Inc.

- **Instance group name**—For example: google.instance-group.myInstanceGroup. To view or create an instance group in the Google console, select **Compute Engine** > **Instance Group**.

User-defined Labels

Panorama uses up to 16 user-defined labels. If you have more than 16 labels, Panorama sorts your user-defined labels alphabetically and uses the first 16 tags.

Review the Google requirements for label key-value pairs: Keys have a minimum length of 1 character and a maximum length of 63 characters, and cannot be empty. Values can be empty, and have a maximum length of 63 characters.

To create or view labels in the GCP console, go to **Compute Engine** > **VM Instances** and select **Show Info Panel**. Select one or more VMs and in the **Info** Panel, select **Labels**. Click **+Add a label**, add a key and value, and click **Save**.

User-defined Network Tags

Panorama uses up to 8 user-defined network tags, If you have more than 8 tags, Panorama sorts your user-defined labels alphabetically and uses the first 8 tags.

Note that Google limits network tags as follows:

- Maximum 63 characters per tag.
- You can use lowercase letters, numbers, and dashes; a tag must start with a lowercase letter, and end with a number or a lowercase letter.

To create or view network tags in the GCP console, go to **Compute Engine** > **VM Instances** and select an instance. **Edit** the instance, and scroll down to **Network Tags**, enter tags (separated by commas), and **Save**. See Configuring Network Tags.

Configure VM Monitoring with the Panorama Plugin for GCP

After you tag your GCP assets and create a service account, make your assets available to Panorama so you can set up VM monitoring.

Prepare Panorama to Configure VM Monitoring

Follow these steps to enable Panorama to manage and monitor your GCP assets. Any VM deployed in GCP can be a managed device in Panorama.

STEP 1 | In Panorama, add the VM-Series firewalls and other VMs associated with your GCP project as managed devices.

STEP 2 | Add a Device Group and assign managed devices to it. A Device Group is a group of firewalls or virtual systems that you want to manage as a group.

A VM can be a member of only one Device Group. Plan your Device Groups carefully.

STEP 3 | Add a template. Name the template and accept the default VPC.

STEP 4 | Add a template stack. **Add** the stack, **Add** the template you just created, and select your devices.

STEP 5 | Commit the changes.

 ©2024 Palo Alto Networks, Inc.

Set Up VM Monitoring

STEP 1 | If you have not done so, Install the Panorama Plugin for GCP.

STEP 2 | Log in to the Panorama web interface and select **Panorama** > **Google Cloud Platform**.

STEP 3 | Set up VM monitoring.

1. Configure general settings.
   1. Select **Panorama** > **Google Cloud Platform** > **Setup** > **General**. To edit the settings, click the gear.
      - Check **Enable Monitoring** to permit VM monitoring on all projects for which you configure a service account.
      - Enter the **Monitoring Interval** in seconds. This is the length of time between tag retrieval events.
2. **Add** a notify group. A notify group is a list of Device Groups to which Panorama pushes IP-address-to-tag mappings and updates.

   A project can have only one notify group.

   1. **Select Panorama** > **Google Cloud Platform** > **Setup** > **Notify Groups** and click **Add**.
   2. Enter a **Name** to identify the group of firewalls to which Panorama pushes the VM information (IP address-to-tag mappings) it retrieves.
   3. Select the **Device Groups** to which Panorama will push the VM information (IP address-to-tag mappings) retrieved from your project. The VM-Series firewalls use the update to determine the current member list for Dynamic Address Groups referenced in Security policy.

      Plan your Device Groups carefully.

   4. Select predefined or custom tags.
      - **Select All 8 Predefined Tags**—Choose this option to select all predefined attributes (tags).
      - **Custom Tags**—Choose this option to create tag lists for predefined attributes, user-defined labels, and user-defined network tags.
   5. - Make sure to include all relevant Device Groups in a single notify group.
      - If you want to deregister the tags that Panorama has pushed to a firewall included in a notify group, you must delete the monitoring definition.
      - To register tags to all virtual systems on a firewall enabled for multiple virtual systems, you must add each virtual system to a separate Device Group on Panorama and assign the Device Groups to the notify group. Panorama will

 ©2024 Palo Alto Networks, Inc.

register tags to only one virtual system, if you assign all the virtual systems to one Device Group.

3. **Add** a GCP Service Account Credential.
   - Name the service account credential.
   - (Optional) Enter a description of the service account.
   - **Browse** to upload the JSON file generated when you created the service account.

   You must use the Panorama web interface. You cannot use the CLI to add a service account

   You can only use a service account for one credential. Do not create multiple credentials from a single JSON file.

After you add a service account credential, you can validate the credential from your Panorama command line:

```
request plugins gcp validate-service-account <svc-acct-credential-name>
```

STEP 4 | Create a **Monitoring Definition**.

A monitoring definition consists of the service account credential for your project and a notify group. All the networking assets in your project are monitored, and the tags retrieved are pushed to the Device Groups you list in your monitoring definition. When you add a new monitoring definition, it is enabled by default.

A project can have only one monitoring definition, and a monitoring definition can include only one notify group.

1. **Select Panorama** > **Google Cloud Platform** > **Monitoring Definition** and click Add.
2. **Name** the monitoring definition.
3. Enter an optional **Description** for the project and assets you are monitoring.
4. Select the **Service Account** credential you created in the previous step.
5. Select a **Notify Group**.
6. **Enable** monitoring for the elements associated with this service account.

STEP 5 | **Commit** the changes on Panorama.

Verify that the status for the Monitoring Definition displays as Success. If it fails, verify that you entered the project ID accurately and provided the correct keys and IDs for the service.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Verify that you can view the VM information on Panorama, and define the match criteria for Dynamic Address Groups.

On HA failover, the newly active Panorama attempts to reconnect to Google Cloud Platform and retrieve tags for all monitoring definitions. If there is an error with reconnecting even one monitoring definition, Panorama generates a system log message:

```
Unable to process subscriptions after HA switch-over; user-
intervention required.
```

If you see this error, fix the issue in Panorama. For example, remove an invalid subscription or provide valid credentials, and commit your changes to enable Panorama to reconnect and retrieve the tags for all monitoring definitions.

Even when Panorama is disconnected from Google Cloud Platform, the firewalls have the list of all tags that had been retrieved before failover, and can continue to enforce policy on that list of IP addresses. When you delete a monitoring definition, Panorama removes all tags associated with registered VMs. As a best practice, configure action-oriented log forwarding to an HTTPS destination *from Panorama so that you can take immediate action.*

 ©2024 Palo Alto Networks, Inc.

Auto Scaling the VM-Series Firewall on Google Cloud Platform

The Panorama plugin for Google Cloud Platform (GCP) version 2.0.0 assists you in deploying the VM-Series firewall in GCP and enables Panorama to manage VM-Series firewalls securing VM monitoring or auto scaling deployments in GCP. Using Panorama for centralized policy and firewall management increases operational efficiency in managing and maintaining a distributed network of firewalls.

With Panorama maintaining your GCP managed instance groups you can create application enablement policies that protect and control the network.

The auto scaling deployment supports using a shared VPC network configuration or VPC network peering to create a common VPC network in which a host project contains shared VPC networks and the VM-Series firewalls, and a service project contains a vm-based or container-based application deployment (a Kubernetes cluster). Palo Alto networks supplies templates to help you deploy the VM-Series firewalls in the host project and deploy an optional sample application in the service project.

BYOL and PAYG licenses can be used for the VM-Series firewalls. During licensing, VM-Series firewall instances talk directly to the Palo Alto Networks license server.

If you choose BYOL your deployment can deactivate license instances in response to a scale-down event. If a VM-Series firewall's deployment information is configured in the Panorama plugin for GCP and the firewall is automatically removed, Panorama detects the firewall status and automatically deregisters the firewall.

- Auto Scaling Components for Google Cloud Platform
- Deploy GCP Auto Scaling Templates

Auto Scaling Components for Google Cloud Platform

Typical GCP auto scaling deployments use a host project and a service project and form a common VPC network between the two. The Panorama plugin for GCP can secure an auto scaling deployment in a single project with host and service VPCs, or host and service projects in a shared VPC or peered VPC network configuration, where the host project contains the VM-Series firewalls and the shared VPC networks, and the service project contains your application deployment. If your application is deployed in a Kubernetes cluster, a peered VPC is required.

- Auto Scaling Requirements
- Prepare to Deploy the Auto Scaling Templates

Auto Scaling Requirements

- **General Requirements**—Ensure your environment meets the basic requirements.
- **Panorama Plugin for GCP**—If you have not done so, Install the Panorama Plugin for GCP.

> *If you previously installed the Panorama plugin for GCP version 1.0.0, remove it before you install 2.0.X. You cannot upgrade.*

 ©2024 Palo Alto Networks, Inc.

- **Palo Alto Networks Auto Scale templates version 1.0**—Palo Alto Networks provides the templates to deploy VM-Series firewall instances in the host project and configure and deploy a sample application in a service project. See About the Auto Scaling Templates for more about the templates.

  Download the templates from GitHub. The zip file contains separate zip files for the firewall and application templates.

Prepare to Deploy the Auto Scaling Templates

Complete the following tasks before you deploy the auto scaling templates.

- Prepare a Host Project and Required Service Accounts
- Obtain a Licensing API Key
- Configure the Panorama Plugin for GCP to Secure an Auto Scaling Deployment
- Prepare a VM-Series Firewall Bootstrap Package for Auto Scaling

Prepare a Host Project and Required Service Accounts

You need a host project and a service project to form the shared VPC topology that supports the firewall and application templates. You can create a new host project or prepare an existing project to act as your host.

To set up the Shared VPC an organization administrator must grant the host project administrator the Shared VPC Admin role. The Shared VPC Admin can enable a project to act as a host, and grant the Service Project Admin role to the service project administrator. Review the GCP documentation on Administrators and IAM roles.

STEP 1 | In the GCP console, create a GCP project to act as the host. If you want to use an existing project, skip to the next step.

To create a new project, select your organization or **No organization**, click **New Project** and fill in your project information. Note, this is your only chance to **EDIT** the project ID.

The Google Cloud SDK must be installed and configured so that you can authenticate with your host project from the CLI. You will use the command line interface to deploy the firewall template and the application template, and to attach the service project to the host project.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enable APIs and services required for auto scaling. The required APIs are:

- Cloud Pub/Sub API
- Cloud Deployment Manager API
- Cloud Storage API
- Compute Engine API
- Google Compute Engine Instance Group Manager API
- Google Compute Engine Instance Group Updater API
- Google Compute Engine Instance Groups API
- Kubernetes Engine API
- Stackdriver API
- Stackdriver Logging API
- Stackdriver Monitoring API

You can enable APIs from the GCP console or the GCP CLI, as shown below.

Enable APIs from the GCP console

1. Select the host project, and from the Navigation menu, select **APIs & Services**.
2. Search for and view each required API.
3. **ENABLE** any APIs that do not display the "API enabled" status.

Enable APIs from the CLI

1. In the CLI, view your configuration to ensure that you are in the correct project.

```
gcloud config list
```

   If not, set the project as follows:

```
gcloud config set project <project-name>
```

2. Issue the following commands to enable the required APIs.

```
gcloud services enable pubsub.googleapis.com
gcloud services enable deploymentmanager.googleapis.com
gcloud services enable storage-component.googleapis.com
gcloud services enable compute.googleapis.com
gcloud services enable replicapool.googleapis.com
gcloud services enable replicapoolupdater.googleapis.com
gcloud services enable resourceviews.googleapis.com
gcloud services enable container.googleapis.com
gcloud services enable stackdriver.googleapis.com
gcloud services enable logging.googleapis.com
```

 ©2024 Palo Alto Networks, Inc.

```
gcloud services enable monitoring.googleapis.com
```

3. Confirm that the required APIs are enabled.

```
gcloud services list --enabled
```

STEP 3 | Create a service account for deploying the VM-Series firewall, and assign the IAM roles required for auto scaling a service or a Kubernetes cluster.

When you configure the firewall templates you add the email address for this service account to the VM-Series firewall `.yaml` file. Within the host project, the template uses credentials

 ©2024 Palo Alto Networks, Inc.

from this service account to create a host VPC with subnets, deploy VM-Series firewalls in the VPC, configure Stackdriver custom metrics, create a Pub/Sub topic, and more.

1. In the GCP console select **IAM & Admin** > **Service accounts** and select **+CREATE SERVICE ACCOUNT**.

   Fill in the service account details and click **CREATE**.
2. Give the service account permission to auto-scale resources in this project.

   Select a role type from the drop menu, and on the right, select an appropriate access level. For example, select Project > Editor. You can select multiple roles for a service account.

   - Compute Engine > Compute Admin
   - Compute Engine > Compute Network User
   - Pub/Sub > Admin
   - Monitoring > Monitoring Metric Writer
   - Stackdriver > Stackdriver Accounts Editor
   - Storage > Storage Admin
   - (GKE only) Kubernetes > Kubernetes Engine Cluster Admin
   - (GKE only) Kubernetes > Kubernetes Engine Viewer

   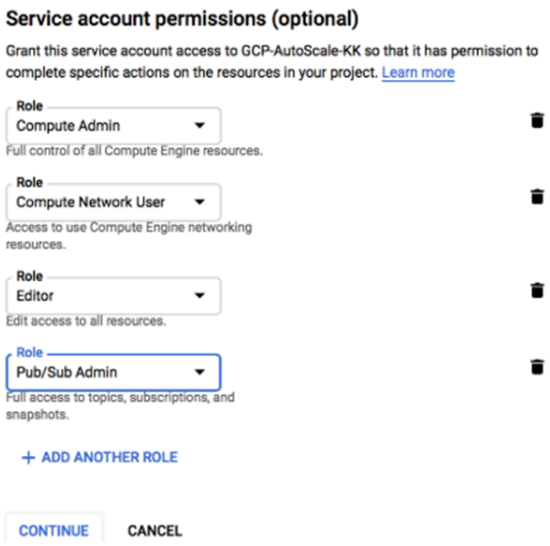

   Continue when you are finished adding roles.
3. Click **+CREATE KEY** to create a key for the host service account.
   - (Optional) Add email addresses to grant other users or administrators access to this service account.
   - Click JSON to download the private key in JSON form.
   - Store the key in a safe location. You will need this key when you Deploy GCP Auto Scaling Templates.
4. Click **DONE**.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Create a service account that a Panorama administrator can use to interact with this host project.

1. In the GCP console select **IAM & Admin** > **Service accounts** and select **+CREATE SERVICE ACCOUNT**.
2. Fill in the service account details and click **CREATE**.
3. Grant service account access.

   Select a role type from the drop menu, and on the right, select an appropriate access level. For example, select Project > Editor. You can select multiple roles for a service account.

   - ❑ Compute Engine > Compute Viewer
   - ❑ Deployment Manager > Viewer
   - ❑ Pub/Sub > Admin

   Click **CONTINUE**.
4. Click **+CREATE KEY** to create a key for the host service account.
   - (Optional) Add email addresses to grant other users or administrators access to this service account.
   - Select JSON to download the private key in JSON form.
   - Store the key in a safe location. You will need this key when you Configure the Panorama Plugin for GCP to Secure an Auto Scaling Deployment.

STEP 5 | (optional) In the CLI, ensure you can communicate with your new host project.

1. Set your project to the host project you just created.

   ```
   gcloud set project <your-autoscale-host-project-name>
   ```
2. Create a configuration for auto scaling. Your new configuration is automatically activated unless you disable activation.

   ```
   gcloud config configurations create <CONFIGURATION_NAME> gcloud config list
   ```

Obtain a Licensing API Key

You need a Licensing API key so Panorama can license and de-license managed assets in GCP.

STEP 1 | Log in to the Support portal and select**Assets** > **Licensing API** and click **Enable**. The key is displayed.

Only a Super User can view the Enable link to generate this key. See How to Enable, Regenerate, Extend the Licensing API Key.

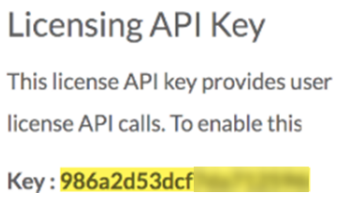

STEP 2 | Select the key and copy it.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | From the CLI, SSH in to Panorama and issue the following command, replacing <key> with the API key you copied from the support portal:

```
request license api-key set key <key>
```

```
API Key is successfully set
```

Configure the Panorama Plugin for GCP to Secure an Auto Scaling Deployment

In Panorama, create assets to support the auto scaling firewall deployment.

STEP 1 | Create a template, and a template stack that includes the template, and **Commit** the changes.

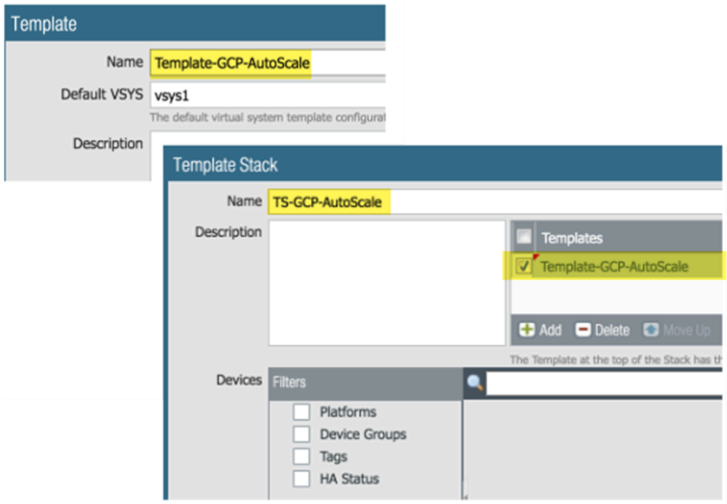

STEP 2 | In the **Network** context, select either the template or the template stack. Select **Virtual Routers** and **Add** a virtual router.

When the firewall template creates static routes, they are added to this virtual router.

Define only one router for the auto scale deployment.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | In the **Network** context, select the template you created, select **Interfaces** and **Add Interface**.

- On the Config tab, select a slot, select the **Interface name** and select the Layer3 **Interface Type**. From the **Security Zone** menu, select **New Zone**, name the zone Untrust and click **OK**.
- On the **IPv4** tab enable **DHCP Client** and **Automatically create default route pointing to default gateway provided by server** (enabled by default) and click **OK**.

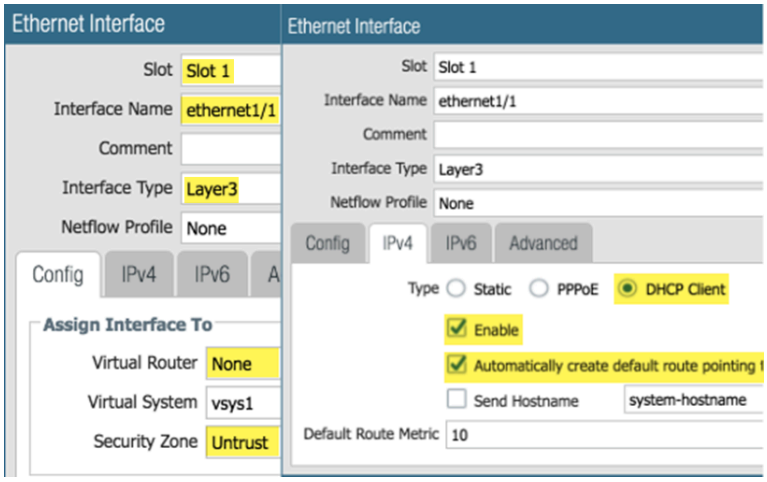

STEP 4 | Add the ethernet1/2 (Trust) Layer 3 interface.

- On the Config tab, chose the same slot as the previous step, select the **Interface name** (ethernet1/2), and select the Layer3 **Interface Type**. From the **Security Zone** menu, select **New Zone** name the zone Trust and click **OK**.
- On the **IPv4** tab enable **DHCP Client**, disable **Automatically create default route pointing to default gateway provided by server** and click **OK**.

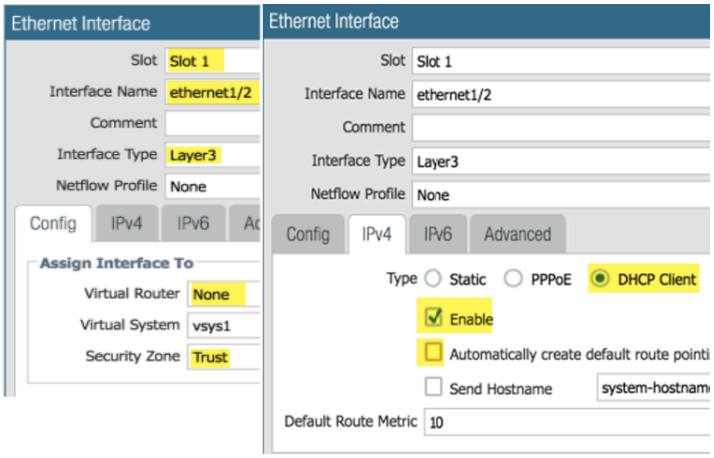

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Return to your template stack and the virtual router you created earlier. Place the untrust and trust interfaces (ethernet1/1 and ethernet1/2) in the virtual router, and click **OK**.

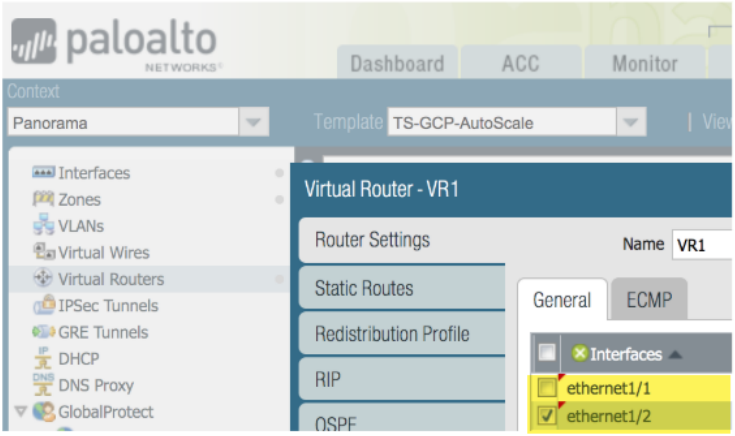

STEP 6 | Configure Stackdriver for your auto scaling deployment.

You must have the VM-Series plugin on Panorama to configure Stackdriver.

1. In the **Device** context, select the template stack you created earlier from the Template drop menu.
2. Select **Device** > **VM-Series** > **Google** and click the Edit cog ( ). Enable **Publish PAN-OS metrics to Stackdriver**.

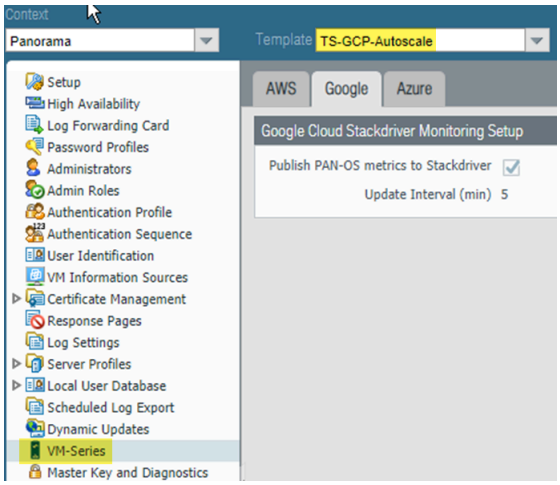

3. Commit your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Create a Device Group that references the template or template stack you created in step 1. This Device Group will contain the VM-Series firewalls you create with the firewall template.

1. Add a security policy that allows web-browsing traffic from Untrust to Trust.

   In the Policies context, select the Device Group you just created. Select **Security** > **Pre Rules** and **Add** the following security policy.

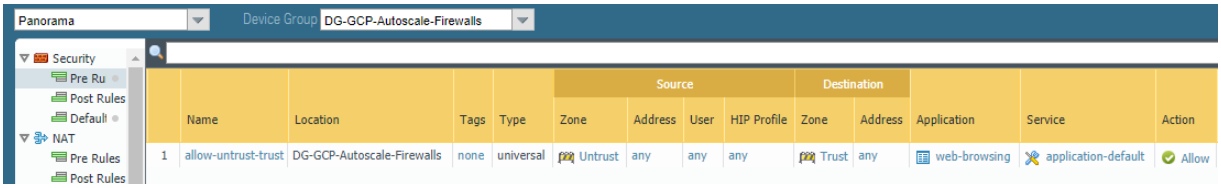

STEP 8 | Set up the GCP service account for the host project.

1. In the Panorama context, expand Google Cloud Platform, select **Setup**, and click **Add**.
2. Supply a name and description for the host service account you created in Step 4.
3. Upload the JSON credentials file you created in Step 4.4.

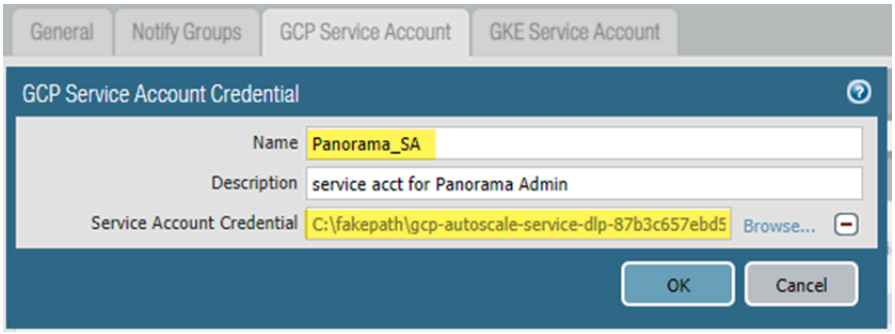

After you add a service account credential, you can validate the credential from your Panorama command line (you cannot validate from the web interface):

```
request plugins gcp validate-service-account
 gcp_service_account <svc-acct-credential-name>
```

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Set up auto scaling on the Panorama plugin for GCP.

1. In the Panorama context, expand Google Cloud Platform, select AutoScaling, and click **Add**.
2. Supply the Firewall Deployment Name and an optional description for the deployment.
3. For the GCP Service Account Credential, supply the GCP service account name from Step 8.

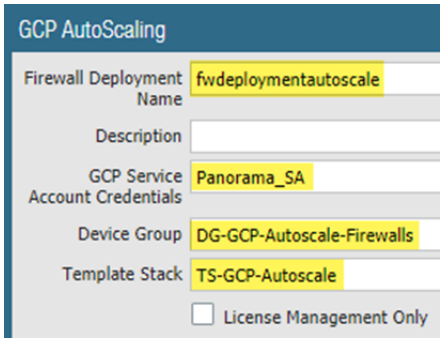

4. Chose the Device Group you created in Step 7, and the Template Stack you created in Step 1.
5. Disable **License Management Only** to ensure traffic is secured.

STEP 10 | Commit your changes.

Prepare a VM-Series Firewall Bootstrap Package for Auto Scaling

During bootstrap, the initial request from the firewall provides the host IP address and serial number, and the VM auth key so Panorama can validate the VM auth key and add the firewall as a managed device. Panorama can then assign the firewall to the appropriate device group and template so that you can centrally configure and administer the firewall using Panorama.

In this case, you must generate a VM auth key on Panorama and include the key in the init-cfg.txt file that you use for bootstrapping. The VM auth key allows Panorama to authenticate the newly bootstrapped VM-Series firewall. The bootstrap package must include.

- In the /config directory, an init-cfg.txt file that includes the Panorama IP address
- In the /license directory, the VM authentication key in a file named authcodes.

  The lifetime of the key can vary between 1 hour and 8760 hours (1 year). After the specified time, the key expires and Panorama will not register VM-Series firewalls without a valid auth-key in this connection request.

STEP 1 | Set up a Google storage bucket with the folders required to Bootstrap the VM-Series Firewall on Google Cloud Platform. You can use an existing bootstrap package or create a new bootstrap package, for these folders.

STEP 2 | Edit the values in the sample `init-cfg.txt` file to customize the file for your environment.

The firewall templates include a sample `init-cfg.txt` file.

Parameter	Value	Comment
type	dhcp-client	

 ©2024 Palo Alto Networks, Inc.

Parameter	Value	Comment
hostname	<pa-vm>	Optional name you assigned when you prepared the host project. Only required if a specific host is necessary, and dhcp-send-hostname is no.
vm-auth-key	<vmauthkey>	A key that Panorama must validate before adding a firewall as a managed device. See Generate the VM Auth Key On Panorama.
panorama-server	<panorama-ip>	The IP address of the Panorama management device you configured in Configure the Panorama Plugin for GCP to Secure an Auto Scaling Deployment
tplname	<template-stack-name>	The template stack you created in Configure the Panorama Plugin for GCP to Secure an Auto Scaling Deployment.
dgname	<dg-name>	The name of the Device Group you created in the Panorama Plugin for GCP.
dns-primary		Your primary DNS server.
dns-secondary		Your secondary DNS server.
dhcp-send-hostname	yes	Leave as is.
dhcp-send-client-id	yes	Leave as is.
dhcp-accept-server-hostname	yes	Leave as is.
dhcp-accept-server-domain	yes	Leave as is.

STEP 3 | Upload your edited `init-cfg.txt` file to the `/config` folder in your bootstrap package.

STEP 4 | If you are using BYOL, create a text file named `authcodes` (no extension), add your auth code, and upload the file to the `/license` folder.

 ©2024 Palo Alto Networks, Inc.

Deploy GCP Auto Scaling Templates

- About the Auto Scaling Templates
- Deploy the Firewall Template
- Prepare a Service Project
- Configure the Shared VPC
- Deploy the Application Template
- Onboard a New Application
- Sample GKE Service Templates

About the Auto Scaling Templates

Download the Palo Alto Networks auto scaling templates from https://github.com/PaloAltoNetworks/GCP-AutoScaling. The zip file contains separate zips for firewall and application templates. Each zip is a template directory containing several files, but you only need to edit the YAML files.

- Firewall Templates
- Application Template

Firewall Templates

The firewall directory files create VM-Series firewalls and other deployment resources. They create new networks and the familiar subnetworks for the VM-Series firewall: management, untrust, and trust. They also deploy a Cloud Pub/Sub messaging service to relay information from GCP to the Panorama plugin for GCP. With this infrastructure in place, the plugin can leverage dynamic address groups to apply security policy on inbound traffic routed to services running on GCP, and use auto scale metrics to deploy VM-Series firewalls to meet increased demand for application workload resources or to eliminate firewalls that are no longer needed.

To configure your load balancer, edit the `.yaml` file for an external application load balancer (ALB) or network load balancer (NLB).

- **ALB** (HTTP External Load Balancer)

  To customize an ALB, edit `vm-series-fw-alb.yaml`.

  HTTP external load balancer is a proxy-based load balancer that performs SNAT and DNAT on the inbound traffic from Internet. The HTTP load balancer is designed to support only the 80 and 8080 TCP ports.

  To support multiple applications using HTTP load balancer in load balancer sandwich architecture, we can use the GCP HTTP load balancer *urlMap* and *namedPort* to map different URLs to different ports in the load balancer. In turn, the VM-Series firewall can translate the ports to different applications, each represented by one internal load balancer per application.

 ©2024 Palo Alto Networks, Inc.

- **NLB** (TCP Load Balancer)

  To customize an NLB, edit `vm-series-fw-nlb.yaml`.

  TCP load balancer is a non-proxy based load balancer, which means it doesn't perform NATing on inbound traffic from the Internet.

  TCP load balancer in GCP allows adding multiple frontend IP addresses with an arbitrary port, making it possible to support multiple applications.

  Another advantage of TCP load balancer is that the original client IP address is preserved, which is desirable for some applications.

Application Template

The application directory provides a sample application. You configure and deploy an internal load balancer (ILB) to enable your application servers to subscribe to the Pub/Sub service and communicate with your VM-Series firewalls and the GCP plugin on Panorama.

To customize the application template, edit `apps.yaml` as described in Deploy the Firewall Template and Application Template.

Deploy the Firewall Template

Edit the Firewall Templates from the host project.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Edit the `vm-series-fw-nlb.yaml` or `vm-series-fw-alb.yaml` environment variables to reflect your cloud environment.

The sample in this workflow is for the NLB. See vm-series-fw-nlb.yaml and vm-series-fw-alb.yaml for further explanation of the template parameters.

```
properties:
    region: us-east1
    zones:
    -us-east1-b
    # Do not modify the lb-type field.
    lb-type: nlb
    cloud-nat: yes
    forwarding-rule-port: 80
```

```
# Only one app is allowed
    urlPath-namedPort-maps:
    - appName: app1
```

```
# ssh key PUBLIC:
```

 ©2024 Palo Alto Networks, Inc.

```
    - optional
```

The autoscaling firewall template requires you to enter the value in single quotes and prepend the key with **admin:** followed by a space. This is the same convention used for the Google Marketplace template, as detailed in SSH Key Pair. For example:

```
    sshkey: 'admin: ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAABAQDe0gJHd8okxPGWXsmdTdcZBJNI6ONT/NSz6INs2CNtKW
[redacted]
oTKXL8t0SRnOKaKV73NR5KnfpsNfwGxG8aQtEkeMCZIxX+6WOnRf/N4K3
yourname@[redacted].paloaltonetworks.local'
```

```
    bootstrap-bucket: bootstrap-autoscale
```

```
    image: vmseries-byol-814
    machine-type: n1-standard-4
```

For the service-account, supply the email address for the host project service account you created earlier (step 3).

```
     service-account: sa-pan@gcp-autoscale-
kk.iam.gserviceaccount.com
```

The fw-instance-tag value will be the managed instance group name in the deployment.

```
     fw-instance-tag: vm-series-fw
```

Choose one metric for auto scaling. Possible values are: panSessionActive, panSessionUtilization, DataPlaneCPUUtilizationPct, DataPlanePacketBufferUtilization, or panSessionUtilization.

```
metric: custom.googleapis.com/VMSeries/panSessionActive
```

```
max-size: 2
min-size: 1
target-type: GAUGE
util-target: 100
```

```
# Greenfield deployment
mgmt-network-cidr: 172.22.2.0/24
untrust-network-cidr: 172.22.1.0/24
trust-network-cidr: 172.22.3.0/24
mgmt-network-access-source-range:
- 199.167.54.229/32
- 199.167.52.5/32
mgmt-network-access-ports:
```

 ©2024 Palo Alto Networks, Inc.

```
- 22
- 443
```

STEP 2 | Deploy the firewall template.

```
gcloud deployment-manager deployments create <your-template>
--config apps.yaml
--automatic-rollback-on-error
```

Take note of the outputs the CLI prints after the deployment—the subnet names, the deployment name, and the Panorama Pub/Sub topic name. You need these values to configure the Shared VPC and for the application template deployment.

The firewall deployment name must be configured in the Panorama plugin for GCP auto scaling definition.

Prepare a Service Project

Create a separate service project, or choose an existing project, for your application.

To learn more about host and service projects in a shared VPC, see the Shared VPC Overview, and review the Administrators and IAM roles. A host project administrator must have the proper role to set up the Shared VPC and make the application project a service project for the host project. See the instructions in Provisioning Shared VPC.

STEP 1 | Enable the service project APIs from the GCP console or the CLI.

The required APIs are:

- Cloud Deployment Manager API
- Cloud Pub/Sub API
- Compute Engine API

Enable APIs from the GCP console

1. Select the service project, and from the Navigation menu, select **APIs & Services**.
2. Search for and view each required API.
3. **ENABLE** any APIs that do not display the “API enabled” status.

Enable APIs from the CLI

1. In the CLI, view your configuration to ensure that you are in the correct project.

   ```
   gcloud config list
   ```

   If not, set the project as follows:

   ```
   gcloud config set project <project-name>
   ```

2. Issue the following commands to enable the required APIs.

   ```
   gcloud services enable deploymentmanager.googleapis.com
   ```

 ©2024 Palo Alto Networks, Inc.

```
gcloud services enable pubsub.googleapis.com
gcloud services enable compute.googleapis.com
```

3. Confirm that the required APIs are enabled.

```
gcloud services list --enabled
```

STEP 2 | Make the application project a service project for the host project.

Add the service account from Service/application project administrator as a member in host project with following roles:

- Compute Network User
- Pub/Sub Admin

STEP 3 | Choose a VPC configuration.

- If the Service project will share the networks in the host project, continue to Configure the Shared VPC.
- If the Service project has its own VPC network for the application deployment, continue to Configure a Peered VPC.

Configure the Shared VPC

After the firewall template is deployed in the host project, configure the service project that supports your applications. An administrator with shared VPC credentials performs these tasks from the host project. To understand more about the host project and service projects in the context of shared VPC, see the Shared VPC Overview.

STEP 1 | Create a shared VPC using the Trust VPC created when you deployed the firewall template.

Set up a shared VPC for the host (firewall) project:

```
gcloud compute shared-vpc enable HOST_PROJECT_ID
```

STEP 2 | Attach the service/application project to the host project.

```
gcloud compute shared-vpc associated-projects add
[SERVICE_PROJECT_ID]--host-project [HOST_PROJECT_ID]
```

Additional options are available to share only specific subnets, rather than all subnets in the host project.

STEP 3 | If you want to use the sample application template to deploy an application, continue to Deploy the Application Template.

If you have already deployed an application and you want to secure it in your auto scaling deployment, go to Manually Onboard an Application to an Existing Auto Scaling Deployment.

If you have deployed a service in a GKE cluster, continue to Onboard a GKE Cluster in a Shared VPC.

 ©2024 Palo Alto Networks, Inc.

Configure a Peered VPC

A VPC network peering connection must be made between two VPCs. If the VPCs are in two different projects, a connection must be created in **both** projects.

STEP 1 | In the host project, peer the Trust VPC network of the Firewall deployment with the Application VPC.

```
gcloud beta compute networks peerings create [PEERING-NAME] \
    --network=[MY-LOCAL-NETWORK] \
    --peer-project [SERVICE-PROJECT-ID] \
    --peer-network [PEER-NETWORK-NAME] \
   [--import-custom-routes] \
   [--export-custom-routes]
```

STEP 2 | In the service project, peer the Trust VPC network of the application deployment with the Trust VPC network of the Firewall deployment.

```
gcloud beta compute networks peerings create [PEERING-NAME] \
    --network=[MY-LOCAL-NETWORK] \
    --peer-project [HOST-PROJECT-ID] \
    --peer-network [PEER-NETWORK-NAME] \
   [--import-custom-routes] \
    [--export-custom-routes]
```

STEP 3 | If you want to use the sample application template to deploy an application, continue to Deploy the Application Template.

If you have already deployed an application and you want to secure it in your auto scaling deployment, go to Manually Onboard an Application to an Existing Auto Scaling Deployment.

If you have deployed a service in a GKE cluster, continue to Onboard a GKE Cluster in a Peered VPC.

Deploy the Application Template

The Service project administrator deploys the Application Template from the service project.

STEP 1 | Create a separate application project (service project) to deploy the application (see Prepare a Service Project).

STEP 2 | Prepare the `apps.yaml` file as outlined in apps.yaml.

STEP 3 | Deploy a new application with the application template and define a label for the named port.

```
gcloud deployment-manager deployments create <your-template>
--config apps.yaml
--automatic-rollback-on-error
```

Continue to View the Onboarded Application in the Panorama Plugin for GCP.

 ©2024 Palo Alto Networks, Inc.

Onboard a New Application

When you use the Application Template to deploy an application, it takes care of the connection to the host project. You can secure applications you did not deploy with the application template, provided they are deployed in a service project with the capabilities described in Prepare a Service Project.

- Manually Onboard an Application to an Existing Auto Scaling Deployment
- Onboard a GKE Cluster

Manually Onboard an Application to an Existing Auto Scaling Deployment

To secure an application you have deployed using an external load balancer and an auto-scaled VM-Series firewall deployment, follow these steps. For each application you onboard, you must supply the application name, the named ports, and the path.

STEP 1 | Prepare to add a new named port and URL path to the HTTP external load balancer created when you deployed the firewall template.

STEP 2 | Update all instance groups named-ports with an additional service name and port values. The following sample onboards the applications app2 and app3.

```
gcloud compute instance-groups set-named-ports
fw-template2-fw-igm-us-east1-b
--zone us-east1-b
--named-ports=app1:80,app2:81,app3:82

gcloud compute instance-groups set-named-ports
fw-template2-fw-igm-us-east1-c
--zone us-east1-c
--named-ports=app1:80,app2:81,app3:82
```

STEP 3 | Create a new http-health-check.

```
gcloud compute backend-services create fw-template2-backend-app3
--protocol="HTTP"
--port-name=app3
--http-health-checks=fw-template2-healthcheck-app3
--load-balancing-scheme="EXTERNAL"
--global
```

STEP 4 | Create a new backend service with the port-name created earlier on the HTTP external load balancer.

```
gcloud compute backend-services create fw-template2-backend-app3
--protocol="HTTP" --port-name=app3
--http-health-checks=fw-template2-healthcheck-app3 --load-
balancing-scheme="EXTERNAL"
```

 ©2024 Palo Alto Networks, Inc.

```
--global
```

Check to see if the new backend service is visible.

```
gcloud compute backend-services list
```

STEP 5 | Edit url-maps and add new path rule. For example:

```
- paths:
    - /app3
    - /app3/*service:
    https://www.googleapis.com/compute/v1/projects/<project-name>/
global/backendServices/fw-template2-backend-app3
```

```
gcloud compute url-maps edit fw-template2-ext-loadbalancer
```

STEP 6 | To secure this application with the VM-Series firewall, manually trigger the pub/sub message through the gcloud CLI. This sends a message to the topic created in the firewall template.

```
gcloud pubsub topics publish
  projects/topics/hj-asg-891ca3-gcp-pavmqa-panorama-apps-
deployment
  --attribute ilb-ip=172.22.9.34,
    app-deployment-name=hj-asg-891ca3-app1,
    ilb-port=80,
    named-port=81,
    network-cidr=172.22.9.0/24,
    fw-deployment-name=hj-asg-891ca3,
    host-project=gcp-pavmqa,
    type=ADD-APP
  --message "ADD-APP"
```

STEP 7 | View the Onboarded Application in the Panorama Plugin for GCP.

STEP 8 | (Optional) To update application attributes, such as ilb-ip, ilb-port, or named-port, issue the pubsub command:

```
gcloud pubsub topics publish projects/gcp-pavmqa/topics/hj-
asg-891ca3-gcp-pavmqa-panorama-apps-deployment
  --attribute ilb-ip=172.22.9.34,
    app-deployment-name=hj-asg-891ca3-app1,
    ilb-port=80,
    named-port=81,
    network-cidr=172.22.9.0/24,
    fw-deployment-name=hj-asg-891ca3,
    host-project=gcp-pavmqa,
    type=UPDATE-APP
  --message "UPDATE-APP"
```

 ©2024 Palo Alto Networks, Inc.

STEP 9 | (Optional) To stop securing the application, issue the following command:

```
gcloud pubsub topics publish projects/gcp-pavmqa/topics/hj-
asg-891ca3-gcp-pavmqa-panorama-apps-deployment
  --attribute ilb-ip=172.22.3.20,app-deployment-name=fw-templ-3-
app-1,
    ilb-port=80,
    named-port=80,
    fw-deployment-name=hj-asg-891ca3,
    type=DEL-APP
  --message "DEL-APP"
```

Onboard a GKE Cluster

To onboard a private GKE cluster, the GCP plugin for Panorama requires the following information.

- In GCP, expose the ELB frontend for the cluster to the GKE service so the VM-Series firewall can get the named port information for the service.
- The cluster API server address.
- The service account credential for the service in which the cluster is deployed, in JSON format.

The GKE cluster name must not exceed 24 characters. This ensures that if you deploy auto scaling in a peered VPC configuration the static route name does not exceed 31 characters.

- Onboard a GKE Cluster in a Shared VPC
- Onboard a GKE Cluster in a Peered VPC
- View the Onboarded Application in the Panorama Plugin for GCP
- View the Deployment Status from the CLI

Onboard a GKE Cluster in a Shared VPC

To onboard the GKE cluster you must share the Host project Trust network VPC with the Service project. See Configure the Shared VPC.

For security reasons, only private clusters should be used in an auto scaling deployment. See Creating a private cluster.

STEP 1 | Set the Host project ID.

```
gcloud config set project [HOST_PROJECT_ID]
```

 ©2024 Palo Alto Networks, Inc.

STEP 2 | (optional) Set compute zone or region for clusters.

If the cluster is zonal, enter the following:

```
gcloud config set compute/zone [COMPUTE_ZONE]
```

If the cluster is regional, enter the following:

```
gcloud config set compute/region [COMPUTE_REGION]
```

STEP 3 | In the Host project, update secondary ranges in the Trust VPC subnet.

```
gcloud compute networks subnets update [TRUST_SUBNETWORK_NAME]
  --add-secondary-ranges
  [PODS_IP_RANGE_NAME] = [POD_RANGE_CIDR],
  [SERVICE_IP_RANGE_NAME]=[SERVICE_RANGE_CIDR]
```

Pods and service IP ranges must be within: 10.0.0.0/8, 172.16.0.0/12 or 192.168.0.0/16, and cannot collide with existing IP ranges in the subnetwork.

STEP 4 | In the Service project, create a private cluster in the shared VPC.

1. Set the Service project ID.

   ```
   gcloud config set project [SERVicE_PROJECT_ID]
   ```

2. Create a private cluster in the shared VPC.

   ```
   gcloud container clusters create  [CLUSTER_NAME]
     --project [SERVICE_PROJECT_ID]
     --zone=[ZONE_NAME]
     --enable-ip-alias
     --enable-private-nodes
     --network projects/[HOST_PROJECT_ID]/global/networks/
   [NETWORK_NAME]
     --subnetwork projects/[HOST_PROJECT_ID]/regions/
   [REGION_NAME]
     /subnetworks/[TRUST_SUBNETWORK_NAME]
     --cluster-secondary-range-name=[PODS_IP_RANGE_NAME]
     --services-secondary-range-name=[SERVICE_IP_RANGE_NAME]
     --master-ipv4-cidr=[MASTER_IPV4_CIDR]
     --enable-master-authorized-networks
     --master-authorized-networks=[PANORAMA_MANAGEMENT_IP/32],
     [MY_MANAGEMENT_IP/32]
   ```

STEP 5 | Check your current cluster context:

```
kubectl config current-context
```

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Check all cluster contexts.

```
kubectl config get-context
```

STEP 7 | Change to another cluster.

```
kubectl config use-context [CONTEXT_NAME]
```

If you created your cluster in the GCP console, generate a kubeconfig entry:

```
gcloud container clusters get-credentials [CLUSTER_NAME]
```

STEP 8 | Create a cluster role in a `.yaml` file—for example, `gke_cluster_role.yaml`.

```
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRole
metadata:
  name: gke-plugin-role
rules:
  - apiGroups:
      - ""
    resources:
      - services
    verbs:
      - list
```

STEP 9 | Apply the cluster role.

```
kubectl apply -f gke_cluster_role.yaml
```

STEP 10 | Create a cluster role binding in a `.yaml` file—for example, `gke_cluster_role_binding.yaml`.

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: gke-plugin-role-binding
subjects:
  - kind: ServiceAccount
    name: [SERVICEACCOUNT_NAME]
    namespace: default
roleRef:
  kind: ClusterRole
  name: gke-plugin-role
  apiGroup: rbac.authorization.k8s.io
```

 ©2024 Palo Alto Networks, Inc.

STEP 11 | Apply the cluster role binding.

```
kubectl apply -f gke_cluster_role_binding.yaml
```

STEP 12 | Create a service account.

```
kubectl create serviceaccount [SERVICEACCOUNT_NAME]
```

STEP 13 | Export the service account secret token in JSON format.

```
MY_TOKEN=`kubectl get serviceaccounts [SERVICEACCOUNT_NAME]
-o jsonpath='{.secrets[0].name}'`

kubectl get secret $MY_TOKEN -o json > [FILE_NAME].json
```

STEP 14 | Get the API server address.

```
kubectl config view --minify | grep server | cut -f 2- -d ":" | tr
 -d " "
```

STEP 15 | In the Panorama plugin for GCP, add the service account information.

Select **Panorama** > **Google Cloud Platform** > **Setup**.

Name the credential, enter a description, and enter the **API server address** from step 14, and for **GKE Service Account Credential**, upload the JSON file you exported in step 13.

After you add a service account credential, you can validate the credential from your Panorama command line (you cannot validate from the web interface):

```
request plugins gcp validate-service-account gke_service_account
 <svc-acct-credential-name>
```

STEP 16 | Set up auto scaling on the Panorama plugin for GCP.

1. In the Panorama context, expand Google Cloud Platform, select AutoScaling, and click **Add**.
2. Supply the **Firewall Deployment Name** and an optional description for the deployment.
3. For the **GCP Service Account Credential**, supply the GCP service account name created in Prepare a Host Project and Required Service Accounts, step 4.
4. Chose the Device Group and the Template Stack you created in when you configured the Panorama plugin.
5. Disable **License Management Only** to ensure traffic is secured.
6. Enter the exact **GKE Cluster Name**.
7. (Optional) Enter a **Description** of the GKE cluster.
8. Enter the **Network CIDR** for the GKE cluster.
9. Select the **GKE Service Account** corresponding to the GKE cluster.

 ©2024 Palo Alto Networks, Inc.

STEP 17 | Commit your changes.

STEP 18 | (Optional) Create and deploy a service template according to Using the Sample GKE Service Templates, or deploy a GKE service in the GCP console. .

Onboard a GKE Cluster in a Peered VPC

To onboard the GKE cluster you must create and peer the Service VPC with the firewall Trust network in the Host project, as described in Configure a Peered VPC.

For security reasons, only private clusters should be used in an auto scaling deployment. See Creating a private cluster.

STEP 1 | Set the project ID.

```
gcloud config set project [PROJECT_ID]
```

STEP 2 | Set compute zone or region for clusters.

If the cluster is zonal, enter the following:

```
gcloud config set compute/zone [COMPUTE_ZONE]
```

If the cluster is regional, enter the following:

```
gcloud config set compute/region [COMPUTE_REGION]
```

STEP 3 | Update the service project VPC network with the secondary IP ranges for the pods and services.

```
gcloud compute networks subnets update [GKE_PEERED_VPC_SUBNETWORK]
--region=[REGION]
--add-secondary-ranges PODS_IP_RANGE_NAME=[ip cidr],
  SERVICE_IP_RANGE_NAME=[ip cidr]
```

STEP 4 | Enable cloud NAT.

Cloud NAT is required to deploy a private cluster.

```
gcloud compute routers create [ROUTER_NAME]
  --network [NETWORK_NAME]
  --region [REGION_NAME]
```

```
gcloud compute routers nats create [NAT_CONFIG_NAME]
  --router-region [REGION_NAME]
  --router [ROUTER_NAME]
  --nat-all-subnet-ip-ranges
```

 ©2024 Palo Alto Networks, Inc.

```
  --auto-allocate-nat-external-ip
```

STEP 5 | Create a new private cluster in the Service VPC.

```
gcloud container clusters create [CLUSTER_NAME]
  --project [SERVICE_PROJECT_ID]
  --zone=[ZONE_NAME]
  --enable-ip-alias
  --network [NETWORK_NAME]
  --subnetwork [SUBNETWORK_NAME]
  --enable-private-nodes
  --cluster-secondary-range-name=[PODS_IP_RANGE_NAME]
  --services-secondary-range-name=[SERVICE_IP_RANGE_NAME]
  --master-ipv4-cidr=[MASTER_IPV4_CIDR]
  --enable-master-authorized-networks
  --master-authorized-networks=[PANORAMA_MANAGEMENT_IP/32],
  [MY_MANAGEMENT_IP/32]
```

STEP 6 | Check your current cluster context:

```
kubectl config current-context
```

STEP 7 | Check all cluster contexts.

```
kubectl config get-context
```

STEP 8 | Change to another cluster.

```
kubectl config use-context [CONTEXT_NAME]
```

If you created your cluster in the GCP console, generate a kubeconfig entry:

```
gcloud container clusters get-credentials [CLUSTER_NAME]
```

STEP 9 | Create a cluster role in a `.yaml` file—for example, `gke_cluster_role.yaml`.

```
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRole
metadata:
  name: gke-plugin-role
rules:
  - apiGroups:
      - ""
    resources:
      - services
    verbs:
      - list
```

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Apply the cluster role.

```
kubectl apply -f gke_cluster_role.yaml
```

STEP 11 | Create a cluster role binding in a `.yaml` file—for example, `gke_cluster_role_binding.yaml`.

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: gke-plugin-role-binding
subjects:
  - kind: ServiceAccount
    name: [SERVICEACCOUNT_NAME]
    namespace: default
roleRef:
  kind: ClusterRole
  name: gke-plugin-role
  apiGroup: rbac.authorization.k8s.io
```

STEP 12 | Apply the cluster role binding.

```
kubectl apply -f gke_cluster_role_binding.yaml
```

STEP 13 | Create a service account.

```
kubectl create serviceaccount [SERVICEACCOUNT_NAME]
```

STEP 14 | Export the service account secret token in JSON format.

```
MY_TOKEN=`kubectl get serviceaccounts [SERVICEACCOUNT_NAME]
-o jsonpath='{.secrets[0].name}'`

kubectl get secret $MY_TOKEN -o json >[FILE_NAME].json
```

STEP 15 | Get the API server address.

```
kubectl config view --minify | grep server | cut -f 2- -d ":" | tr
 -d " "
```

 ©2024 Palo Alto Networks, Inc.

STEP 16 | In the Panorama plugin for GCP, add the service account information.

Select **Panorama** > **Google Cloud Platform** > **Setup**.

Name the credential and enter the **API server address** from Step 15, and upload the JSON file you exported in Step 14.

After you add a service account credential, you can validate the credential from your Panorama command line:

```
request plugins gcp validate-service-account <svc-acct-credential-name>
```

STEP 17 | Set up auto scaling on the Panorama plugin for GCP.

1. In the Panorama context, expand Google Cloud Platform, select AutoScaling, and click **Add**.
2. Supply the **Firewall Deployment Name** and an optional description for the deployment.
3. For the **GCP Service Account Credential**, supply the GCP service account name from Step 16.
4. Chose the Device Group and the Template Stack you created in when you configured the Panorama plugin.
5. Disable **License Management Only** to ensure traffic is secured.
6. Enter the exact **GKE Cluster Name**.
7. (Optional) Enter a **Description** of the GKE cluster.
8. Enter the **Network CIDR** for the GKE cluster.
9. Select the **GKE Service Account** corresponding to the GKE cluster.

STEP 18 | (Optional) In your service project, create and deploy a GKE template according to Using the Sample GKE Service Templates, or deploy a GKE service use the GCP console. Onboard a GKE Cluster

View the Onboarded Application in the Panorama Plugin for GCP

Select **Panorama** > **Google Cloud Platform** > **Autoscaling** to view your onboarded application. The **Details** column is only visible if you have an onboarded application.

Firewall Deployment Name	Project ID	Device Group	Template Stack	Details
gcp-asg-fw-peerbrown0	gcp-pavmqa	GCP_ASG_DG_peerbrown0	GCP_ASG_TS_peerbrown0	Show Status Delicense Inactive VMs Trigger GKE Services Sync
hj-nlb-n642wb	gcp-autoscale-host-250622	gcp-autoscale-dg2	gcp-autoscale-ts2	Show Status Delicense Inactive VMs Trigger GKE Services Sync
hj-asg-891ca3	gcp-pavmqa	gcp-autoscale-dg-891ca3	gcp-autoscale-ts-891ca3	Show Status Delicense Inactive VMs Trigger GKE Services Sync
hj-asg-y892bl	gcp-pavmqa	gcp-autoscale-dg-y892bl	gcp-autoscale-ts-y892bl	Show Status Delicense Inactive VMs Trigger GKE Services Sync

Each link in the Details column triggers an action.

 ©2024 Palo Alto Networks, Inc.

- **Show Status**— View the details for applications onboarded to a GCP VM-Series firewall deployment.

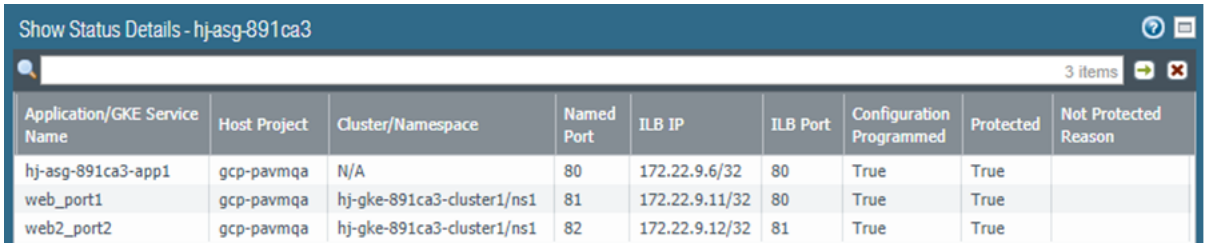

Application/GKE Service Name	Host Project	Cluster/Namespace	Named Port	ILB IP	ILB Port	Configuration Programmed	Protected	Not Protected Reason
hj-asg-891ca3-app1	gcp-pavmqa	N/A	80	172.22.9.6/32	80	True	True	
web_port1	gcp-pavmqa	hj-gke-891ca3-cluster1/ns1	81	172.22.9.11/32	80	True	True	
web2_port2	gcp-pavmqa	hj-gke-891ca3-cluster1/ns1	82	172.22.9.12/32	81	True	True	

  The following fields display information obtained from the selected deployment. You specified these values in the pub/sub message or through GKE cluster service polling.

  - **Application/GKE Service Name**—An application deployment name, or the name of a GKE service.
  - **Host Project**—The name of the host project.
  - **Cluster/Namespace**—A GKE cluster name followed by the namespace for example, **`mycluster/namespace9`**.
  - **Named Port**—The port assigned to the named port for the service.
  - **ILB IP**—The ILB IP address.
  - **ILB Port**—The ILB port number.

    For autoscaling an application, this property is **`ilb-port`** in `apps.yaml`.

    For securing a GKE cluster, this value is the port number of the GKE cluster, as specified in the `.yaml` file you used to deploy the service in your cluster.
  - **Configuration Programmed**— True if a NAT Rule exists, False if not.
  - **Protected**— True when an application is onboarded successfully, or False if onboarding failed. If False, see the **Not Protected Reason** column for an explanation.
  - **Not Protected Reason—** If **Protected** is False, displays the reason the application is not protected. Some common reasons are:
    - **Configuration Programmed**—True if a NAT Rule exists, False if not.
    - **Protected**—True when an application is onboarded successfully, or False if onboarding failed. If False, see the **Not Protected Reason** column for an explanation.
    - **Not Protected Reason**—If **Protected** is False, displays the reason the application is not protected. Some common reasons are:
      - You deployed a UDP service in the GKE cluster.
      - You specified a named port that is already in use. Only one application can listen on a specific named port.
      - You chose the **License management only** option, so we do not program the configuration.
      - No matching label found for GKE services.
- **Delicense Inactive VMs**—Answer **Yes** to trigger the delicensing function for inactive VMs.

 ©2024 Palo Alto Networks, Inc.

- **Trigger GKE Services Sync**—Answer **Yes** to poll the services running in the clusters, and program the NAT, address, and service objects, and static routes if necessary. By default, Panorama automatically polls 10 minutes after the completion of the previous poll.

View the Deployment Status from the CLI

You can use the Panorama CLI to manage deployed applications. The command line actions parallel those described in View the Onboarded Application in the Panorama Plugin for GCP. In the following commands, the **autoscaling_name** is the Firewall Deployment Name you entered in the auto scaling configuration.

- List the onboarded (protected) applications.

```
show plugins gcp show-protected-apps autoscaling_name <fw-deployment-name>
```

- Trigger the delicensing function for firewalls in the specified deployment.

```
request plugins gcp force-delicensing autoscaling_name <fw-deployment-name>
```

- For a GKE deployment, force the plugin to read the pub-sub messages, and sync NAT rules that are programmed based on the pub-sub messages.

```
request plugins gcp gke-service-discovery autoscaling_name <fw-deployment-name>
```

Parameters in the Auto Scaling Templates for GCP

You can download the template `.zip` file from https://github.com/PaloAltoNetworks/GCP-AutoScaling. The `.zip` file contains directories to support firewall templates for network load balancer and application load balancer configurations, and the application template.

The template YAML files have the following general format:

```
#Copyright and license information
    :
    :
imports:                    <do not change>
    :
    :
resources:
    -name: vm-series-fw     <do not change>
    -type:vm-series-fw.py   <do not change
    -properties:
    :
    :
outputs:                    <do not change>
    :
    :
```

In all `.yaml` files, you customize the `resources` properties for your deployment. Do not change anything in the `imports` or `outputs` sections.

 ©2024 Palo Alto Networks, Inc.

- Firewall Templates
- Application Template

Firewall Templates

The following sections detail the parameters for the NLB and ALB `.yaml` files.

- vm-series-fw-nlb.yaml
- vm-series-fw-alb.yaml

vm-series-fw-nlb.yaml

In the `vm-series-fw-nlb.yaml` template, edit the `-properties`.

Parameter	Sample Value	Comment
region	us-central1	https://cloud.google.com/compute/docs/regions-zones
zones - <list of zones>	zones- us-central1-a	If applicable, list multiple zones as follows: zones- us-central1-a- us-central1-b- us-central1-c- us-central1-f
lb-type	nlb	Do not change.
cloud-nat	yes	Do not change.
forwarding-rule-port	80	80 or 8080
urlPath-namedPort-maps-appname	`urlPath-namedPort-maps -MyApplication`	Enter your application name.
sshkey	'admin:ssh-rsa <PASTE KEY>'	Review SSH Key Pair. In single quotes, type **`admin:`** followed by a space, and paste in your key. This is the same convention used for the Google Marketplace template.
bootstrap-bucket	bootstrap-autoscale	The name of the GCP bucket that contains your bootstrap file.
image	vm-series-byol-814	The BYOL image currently available from the Google marketplace.

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
		If you are using PAYG or another license model, the image might be different.
machine-type	n1-standard-4	n1-standard-4 is default for BYOL. If your license permits it, you can use any machine type in Minimum System Requirements for the VM-Series Firewall.
service-account		The unique service account name for the host project.
fw-instance-tag	vm-series-fw	The instance tag you provided in GCP.

metric	custom.googleapis.com/VMSeries/panSessionActive	The custom API path for VM-Series, and your chosen auto scaling metric.
		Supply only one of the following metrics.
		`panSessionActive`
		`panSessionUtilization`
		`DataPlaneCPUUtilizationPct`
		`DataPlanePacketBufferUtilization`
		`panSessionUtilization`

max-size	2	
min-size	1	
target-type	GAUGE	Currently GAUGE is the only valid type.
util-target	100	

To deploy the VM-Series firewall you need a dedicated network and subnetwork for the firewall's managment, untrust, and trust interfaces. Fill out the information for either a greenfield deployment (configure the template to create new networks) or brownfield

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
deployment (use existing networks). Be sure to remove or comment out the network deployment parameters you are not using.		
Greenfield Deployment: Enter values to create management, untrust, and trust networks and subnetworks for the firewall.		
mgmt-network-cidr	172.22.2.0/24	
untrust-network-cidr	172.22.1.0/24	
trust-network-cidr	172.22.3.0/24	

mgmt-network-access-source-range- <permitted-ip-range>	`mgmt-network-access-source-range`
	`- <permitted-ip-range-1>`
	`- <permitted-ip-range-2>`	

mgmt-network-access-ports-<port-number>	`mgmt-network-access`
	`-ports`
	`- 22`
	`- 443`	

Brownfield Deployment: Enter the name of each existing network or subnetwork		
mgmt-network	my-mgmt-network	
mgmt-subnet	my-mgmt-subnet	
trust-network	my-trust-network	
trust-subnet	my-trust-subnet	
untrust-network	my-untrust-network	
untrust-subnet	my-untrust-subnet	

vm-series-fw-alb.yaml

In the `vm-series-fw-alb.yaml` template, edit the `-properties`.

Parameter	Sample Value	Comment
region	us-central1	https://cloud.google.com/compute/docs/regions-zones
zones	zones- us-central1-a	If applicable, list multiple zones as follows:

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
- <list of zones>		zones- us-central1-a- us-central1-b- us-central1-c- us-central1-f
lb-type	alb	Do not change.
cloud-nat	yes	Do not change.
forwarding-rule-port	80	80
connection-draining-timeout	300	The timeout value in seconds.

`urlPath-namedPort-maps:`
`- appname:`
`  namedPort:`
`  urlMapPaths:`
`  - '/app1'`
`  - '/app1/*'`
	`urlPath-namedPort-maps:`
	`- appName: app1`
	`  namedPort: 80`
	`urlMapPaths:`
	`  - '/app1'`
	`  - '/app1/*'`
	`- appName: app2`
	`  namedPort: 81`
	`  urlMapPaths:`
	`  - '/app2'`
	`  - '/app2/*'`
		List your apps and the corresponding named port

sshkey	'admin:ssh-rsa <PASTE KEY>'	Review SSH Key Pair. In single quotes, type **`admin:`** followed by a space, and paste in your key. This is the same convention used for the Google Marketplace template.
bootstrap-bucket	bootstrap-bucket-name	The name of the GCP bucket that contains your bootstrap file.
image	vm-series-byol-814	The BYOL image currently available from the Google marketplace. If you are using PAYG or another license model, the image might be different
machine-type	n1-standard-4	n1-standard-4 is default for BYOL. If your license permits it, you can use any machine

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
		type in Minimum System Requirements for the VM-Series Firewall.
service-account	The unique service account name for the service project.	
fw-instance-tag	vm-series-fw	The instance tag you provided in GCP.

metric	custom.googleapis.com/VMSeries/panSessionActive	The custom API path for VM-Series, and your chosen auto scaling metric.
		Supply only one of the following metrics.
		`panSessionActive`
		`panSessionUtilization`
		`DataPlaneCPUUtilizationPct`
		`DataPlanePacketBufferUtilization`
		`panSessionUtilization`

max-size	2	
min-size	1	
target-type	GAUGE	Currently GAUGE is the only valid type.
util-target	100	Enter the goal utilization target value for the auto scaling.
Greenfield Deployment: Enter values to create management, untrust, and trust networks and subnetworks for the firewall.		
mgmt-network-cidr	192.168.12.0/24	
untrust-network-cidr	192.168.11.0/24	
trust-network-cidr	192.168.11.0/24	
mgmt-network-access-source-range- <permitted-ip-range>	`mgmt-network-access-source-range- <permitted-ip-range-1>- <permitted-ip-range-2>`	

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
mgmt-network-access-ports- <port-number>	`mgmt-network-access-ports- 22- 443`	
Brownfield Deployment: Enter the name of each existing network or subnetwork		
mgmt-network	existing-vpc-mgmt	
mgmt-subnet	existing-subnet-mgmt	
trust-network	existing-vpc-trust	
trust-subnet	existing-subnet-trust	
untrust-network	existing-vpc-untrust	
untrust-subnet	existing-subnet-untrust	

Application Template
apps.yaml

The application template creates the connection between the host project (which contains the VM-Series firewalls) and the service project, which contains the application or services that the firewall deployment secures.

Parameter	Sample Value	Comment
host-project	your-host-project-name	The name of the project containing the VM-Series firewall deployment.
fw-deployment-name	my-vm-series-firewall-name	
region	us-central1	https://cloud.google.com/compute/docs/regions-zones

zones
- <list of zones>
	zones- us-central1-a	If applicable, list multiple zones as follows:
		`zones- us-central1-a -`
		`us-central1-b-`
		`us-central1-c-`
		`us-central1-f`

app-machine-type	n1-standard-2	The machine type for the VM running your application or service. If your license permits

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
		it, you can use any machine type in Minimum System Requirements for the VM-Series Firewall.
app-instance-tag	web-app-vm	You applied this tag (label) in GCP.
sshkey	'admin:ssh-rsa <PASTE KEY>'	Review SSH Key Pair. In single quotes, type **admin:** followed by a space, and paste in your key. This is the same convention used for the Google Marketplace template.
trust-network	<project-name>/<vpc-network-name>	For a shared VPC, the <project-name> is the host project name. For peered VPCs the <project-name> is the Service project name.
trust-subnet	<project-name>/<subnet-name>	For a shared VPC, the <project-name> is the host project name. For peered VPCs the <project-name> is the Service project name.
trust-subnet-cidr	10.2.0.0/24	For a greenfield deployment, the Host project Trust subnet CIDR (the trust-network-cidr parameter in the firewall template). For a brownfield deployment, the CIDR for the Trust network.
vm-series-fw-template-topic	<pubsub-topic>	Enter the topic name created by the firewall deployment. The application template posts a message to the topic to program the firewall configuration to forward traffic.

 ©2024 Palo Alto Networks, Inc.

Parameter	Sample Value	Comment
ilb-port	80	Enter the port number for your application's internal-load-balancer-port. output.
urlPath-namedPort	83	Enter the port number for the urlPath-namedPort output.

Sample GKE Service Templates

These sample templates demonstrate how to configure a GKE service so it is secured by the VM-Series firewall. For the basics on creating your own cluster services, see Creating a private cluster.

- Using the Sample GKE Service Templates
- gke_cluster_role.yaml
- gke_cluster_role_binding.yaml
- web-deployment.yaml
- web-service.yaml
- web-deployment-v2.yaml
- web-service-v2.yaml
- Multiple Ports in a Service

Using the Sample GKE Service Templates

You can create a service template based on the sample content in the `.yaml` files that follow. Typically you create a single `.yaml` file.

To be secured by the VM-Series firewall, services in the cluster must be labeled "panw-named-port=<named_port>" as shown in web-service.yaml or web-service-v2.yaml.

1. Deploy a `.yaml` file as follows:

```
kubectl apply -f [FILE_NAME].yaml
```

2. Configure the VPC deployment.
   - In a shared VPC deployment, launch the GKE cluster in the shared VPC as described in Configure the Shared VPC.
   - In a peered VPC deployment, peer the GKE cluster VPC to the host project Trust network. See Configure a Peered VPC.

After a deployment, you can delete all services deployed in the service template `.yaml` file as follows:

```
kubectl delete -f [FILE_NAME].yaml
```

 ©2024 Palo Alto Networks, Inc.

gke_cluster_role.yaml

```
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRole
metadata:
  name: gke-plugin-role
rules:
  - apiGroups:
      - ""
    resources:
      - services
    verbs:
      - list
```

gke_cluster_role_binding.yaml

```
kind: ClusterRoleBinding
apiVersion: rbac.authorization.k8s.io/v1beta1
metadata:
  name: gke-plugin-role-binding
subjects:
  - kind: ServiceAccount
    name: hj-gke-891ca3-cluster1-sa
    namespace: default
roleRef:
  kind: ClusterRole
  name: gke-plugin-role
  apiGroup: rbac.authorization.k8s.io
```

web-deployment.yaml

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: web
  namespace: default
spec:
  selector:
   matchLabels:
    run: web
  template:
   metadata:
    labels:
      run: web
   spec:
    containers:
    - image: gcr.io/google-samples/hello-app:1.0
      imagePullPolicy: IfNotPresent
      name: web
      ports:
      - containerPort: 8080
       protocol:  TCP
```

 ©2024 Palo Alto Networks, Inc.

web-service.yaml

```
apiVersion: v1
kind: Service
metadata:
  name: web
  namespace: default
  annotations:
    cloud.google.com/load-balancer-type: "Internal"
  labels:
    panw-named-port-port1: "80"
spec:
  ports:
  # the port that this service should serve on
  - name: port1
    port: 80
    protocol: TCP
    targetPort: 8080
  selector:
    run: web
  type: LoadBalancer
```

web-deployment-v2.yaml

```
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: web2
  namespace: default
spec:
  selector:
    matchLabels:
      run: web2
  template:
    metadata:
      labels:
        run: web2
    spec:
      containers:
      - image: gcr.io/google-samples/hello-app:2.0
        imagePullPolicy: IfNotPresent
        name: web2
        ports:
        - containerPort: 8080
          protocol: TCP
```

web-service-v2.yaml

```
apiVersion: v1
  kind: Service
  metadata:
    name: web2
    namespace: default
    annotations:
```

 ©2024 Palo Alto Networks, Inc.

```
      cloud.google.com/load-balancer-type: "Internal"
    labels:
      panw-named-port-port2: "81"
  spec:
    ports:
    # the port that this service should serve on
    - name: port2
      port: 81
      protocol: TCP
      targetPort: 8080
    selector:
      run: web2
    type: LoadBalancer
```

Multiple Ports in a Service

For multiple ports in one service, edit labels and map the target port name and number in the format panw-named-port-<*service-spec-port-name*>, as shown in the sample below.

```
apiVersion: v1
kind: Service
metadata:
  name: carts
  annotations:
    cloud.google.com/load-balancer-type: "Internal"
  labels:
    panw-named-port-carts-http: "6082"
    panw-named-port-carts-https: "6083"
  namespace: default
spec:
  type: LoadBalancer
  ports:
  # the port that this service should serve on
  - name: carts-http
    protocol: TCP
    port: 80
    targetPort: 80
  - name: carts-https
    protocol: TCP
    port: 443
    targetPort: 443
  selector:
    name: carts
```

 ©2024 Palo Alto Networks, Inc.

Set Up a VM-Series Firewall on a Cisco ENCS Network

If you have virtualized the traditional appliance-based network infrastructure at your branch or remote office with the Cisco 5400 Series Enterprise Network Compute System (ENCS) appliance, you can use Enterprise NFV Infrastructure Software (NFVIS) to deploy the VM-Series firewall within your Cisco network. The VM-Series firewall serves as a virtual network function (VNF) with next-generation firewall capabilities to safely enable all applications and protect your branch or remote office users and network from threats.

The Cisco Enterprise Network Compute System (ENCS) appliances combine with Cisco Integrated Services Virtual Routers (ISRV) and NFVIS software to support Software-Defined Branch (SD-Branch) network architectures.

- Plan Your Cisco ENCS Deployment
- Prepare the VM-Series Firewall Image for Cisco ENCS
- Deploy the VM-Series Firewall on Cisco ENCS

Plan Your Cisco ENCS Deployment

In your Cisco SD-Branch, deploy the VM-Series Firewall on the Cisco ENCS appliance as a VNF that provides next generation firewall capabilities to secure your applications and users at the branch office. You can deploy the firewall in a virtual wire, Layer 2, or Layer 3 deployment, and in high availability configuration.

To manage the VM-Series firewall, the Panorama appliance can be deployed on premises or in the cloud. The following topology shows the VM-Series firewall at the branch edge.

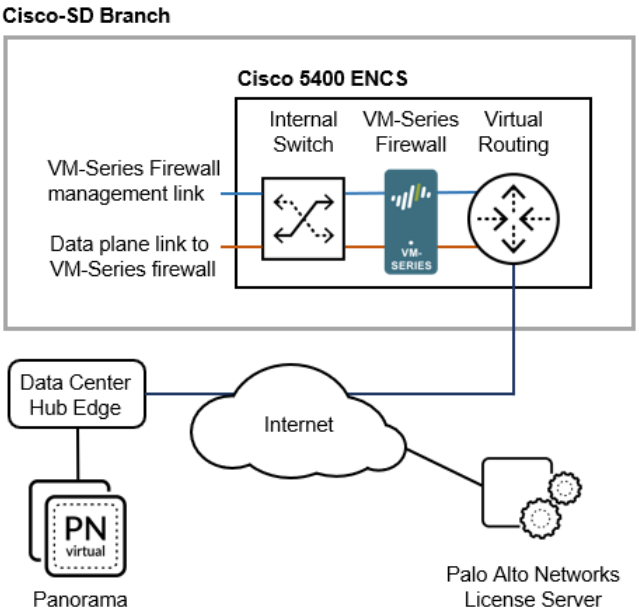

Cisco ENCS Requirements

For supported NFVIS versions and hardware platforms, see the Palo Alto Networks Compatibility Matrix.

- In NFVIS, set up networks and bridges.
  - Create virtual NICs and attach them to a virtual bridge so the ENCS appliance can steer traffic through the VM-Series firewall.

    On the Cisco ENCS appliance, the VM-Series firewall supports up to 8 dataplane interfaces.

    The dataplane interfaces of the VM-Series firewall on Cisco ENCS support Virtio mode only; ENCS SR-IOV and PCI passthrough modes are not supported.
  - Set up network connections for VM-Series firewall management access. If you are using Panorama, ensure that Panorama has network access to manage the firewall you deploy.
- Python 2.7. Required on your local machine if you are using the command line to convert.

VM-Series Firewall and Panorama Requirements

- VM-Series Firewall—The VM-50 and VM-100 are recommended. The VM-300, VM-500, and VM-700 are also supported, provided the ENCS hardware has sufficient resources that can be assigned to the VM-Series firewall. Consult the VM-Series System Requirements to ensure

 ©2024 Palo Alto Networks, Inc.

that the Cisco ENCS appliance has adequate resources to support the VM-Series model you choose.

- qcow2 file (the PAN-OS for VM-Series KVM base image) for PAN-OS 9.1 or later. See Convert a qcow2 File from the Graphical User Interface, Step 2 or Convert a qcow2 File from the Command Line Interface, Step 2.
- VM-Series firewall capacity license and subscription auth codes that meet your requirements. See VM-Series Model License Types. You enter auth codes in the NFVIS user interface, or include the auth codes in the **`authcodes`** text file in the conversion folder as described in Convert a qcow2 File from the Command Line Interface, Step 4.
- With PAN-OS 9.1, the VM-Series firewall on Cisco ENCS supports Virtio with DPDK mode enabled by default.

- Panorama hardware or virtual appliance. While you can deploy a single VM-Series firewall in a Cisco SD-Branch network, it is more common to deploy firewalls in many branches and centrally manage them with Panorama.
  - Panorama version 9.1 or later. The version must be the same or higher than the version on your VM-Series firewall.
  - A VM auth key generated on Panorama. This key allows the VM-Series firewall to authenticate with Panorama.

 ©2024 Palo Alto Networks, Inc.

Prepare the VM-Series Firewall Image for Cisco ENCS

You can convert a PAN-OS qcow2 file from the NFVIS graphical user interface or the command line interface.

- Convert a qcow2 File from the Graphical User Interface
- Convert a qcow2 File from the Command Line Interface

Convert a qcow2 File from the Graphical User Interface

Use the NFVIS graphical user interface to enter image packaging and bootstrap information.

STEP 1 | In NFVIS, go to **VM Life Cycle** > **Image Repository** > **Image Packaging**.

STEP 2 | Fill in the package information as shown below, supplying your own values.

1. Enter a **Package Name** and **VM Version**, and for the **VM Type**, choose **Firewall**.
2. **Enable** the **Serial Console**.
3. Leave the **Sriov Driver(s)** field blank, as SR-IOV is not supported.
4. Select **Local** to choose a qcow2 file you uploaded previously, or click **Upload Raw Images** to upload a qcow2 file.
   - Log in to the Palo Alto Networks Customer Support Portal.

     If you have not already done so, create a support account and register the VM-Series firewall.
   - Select **Support** > **Software Updates** and from the **Filter By** drop-down, select **Pan OS for VM-Series KVM Base Image**, for example, version 9.1.
   - Download the qcow2 image.

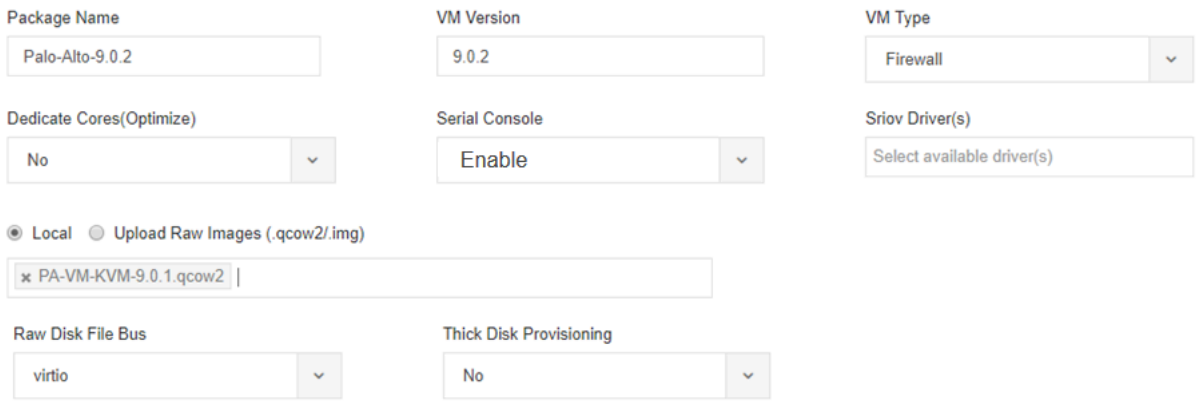

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Upload the bootstrap files.

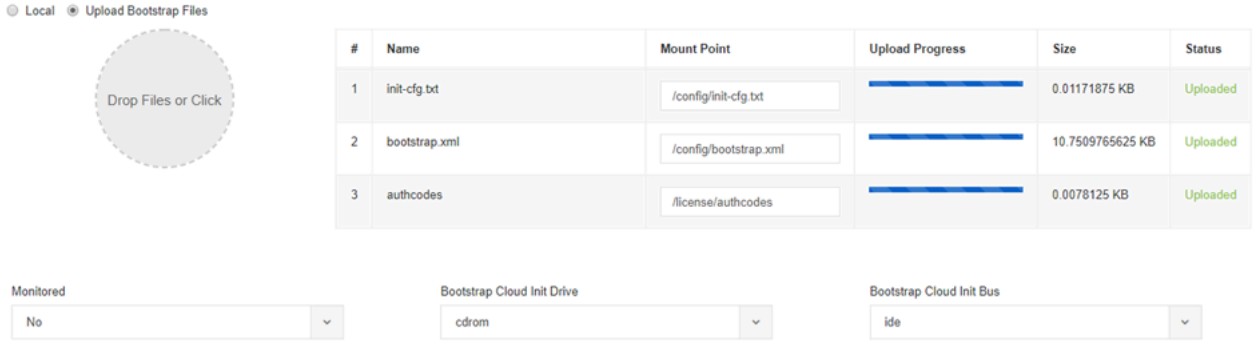

STEP 4 | Set the **Advanced Configuration**.

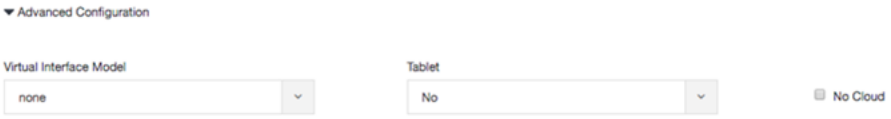

STEP 5 | Enter values for **Custom Properties**.

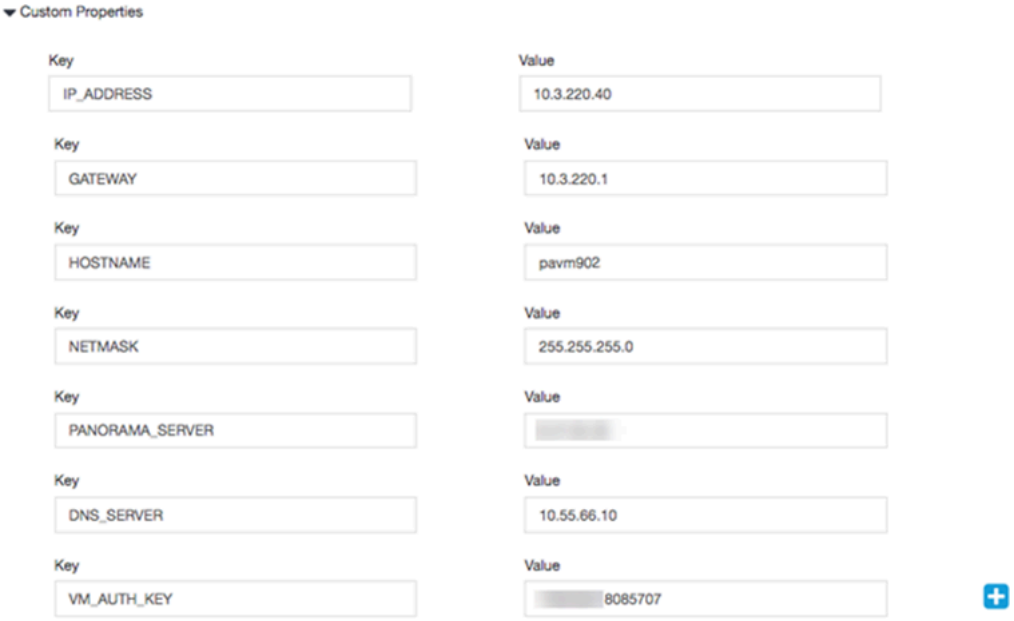

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Set values for your resource requirements and choose the Default profile, or add a profile for the current configuration.

Click **Submit** to save your package.

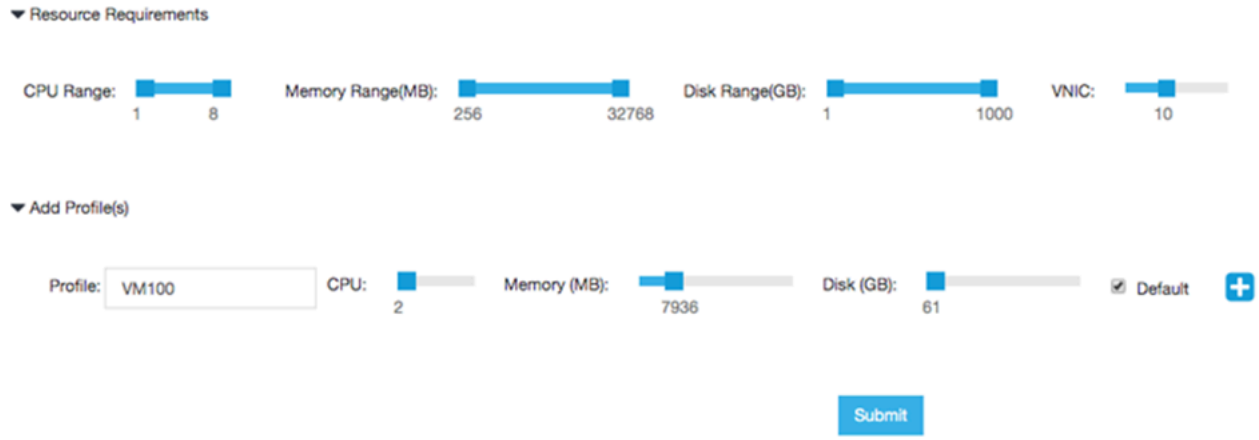

STEP 7 | Click **Register** to register the new image.

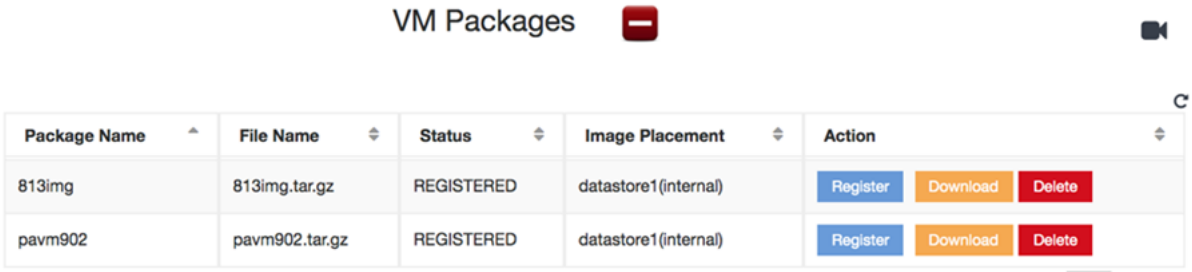

Convert a qcow2 File from the Command Line Interface

To create a bootstrap file from the command line interface, you create the file `image_properties_template.xml` then use the using the VM Image Packaging utility to create a `.tar` file, which you convert using the `nfvpt.py` script. The output is a `tar.gz` file that can be uploaded from the NFVIS user interface.

STEP 1 | Create or choose a folder on your local machine (the conversion folder) in which you want to download and save the files necessary to convert the VM-Series firewall qcow2 image to the Cisco ENCS format.

STEP 2 | Obtain the VM-Series firewall qcow2 image.

1. Log in to the Palo Alto Networks Customer Support Portal.

   If you have not already done so, create a support account and register the VM-Series firewall.
2. Select **Support** > **Software Updates** and from the **Filter By** drop-down, select **Pan OS for VM-Series KVM Base Image**, for example, version 9.1.
3. Download the qcow2 image to the conversion folder.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Create the following init-cfg.txt file in the conversion folder.

```
type=static
ip-address=${IP_ADDRESS}
default-gateway=${GATEWAY}
netmask=${NETMASK}
ipv6-address=
ipv6-default-gateway=
hostname=${HOSTNAME}
vm-auth-key=${VM_AUTH_KEY}
panorama-server=${PANORAMA_SERVER}
panorama-server-2=
tplname=
dgname=
dns-primary=${DNS_SERVER}
dns-secondary=
op-command-modes=jumbo-frame, mgmt-interface-swap**
dhcp-send-hostname=yes
dhcp-send-client-id=yes
dhcp-accept-server-hostname=yes
dhcp-accept-server-domain=yes
```

STEP 4 | Create a text file named `authcodes` (no extension), and enter the auth codes for the VM-Series firewall capacity and subscriptions. Save the file in the conversion folder.

STEP 5 | Create the following `image_properties_template.xml` file in the conversion folder, and supply values for your deployment:

```
<image_properties>
    <vnf_type>FIREWALL</vnf_type>
    <name>pafw</name>
    <version>9.1.0</version>
    <bootup_time>-1</bootup_time>
    <root_file_disk_bus>virtio</root_file_disk_bus>
    <root_image_disk_format>qcow2</root_image_disk_format>
    <vcpu_min>2</vcpu_min>
    <vcpu_max>8</vcpu_max>
    <memory_mb_min>4096</memory_mb_min>
    <memory_mb_max>16384</memory_mb_max>
    <vnic_max>8</vnic_max>
    <root_disk_gb_min>32</root_disk_gb_min>
    <root_disk_gb_max>60</root_disk_gb_max>
    <console_type_serial>true</console_type_serial>
    <sriov_supported>true</sriov_supported>
    <pcie_supported>false</pcie_supported>
    <monitoring_supported>false</monitoring_supported>
    <monitoring_methods>ICMPPing</monitoring_methods>
    <low_latency>true</low_latency>
    <privileged_vm>true</privileged_vm>
    <custom_property>
        <HOSTNAME> </HOSTNAME>
    </custom_property>
    <custom_property>
        <IP_ADDRESS> </IP_ADDRESS>
```

 ©2024 Palo Alto Networks, Inc.

```
    </custom_property>
    <custom_property>
        <NETMASK> </NETMASK>
    </custom_property>
    <custom_property>
        <GATEWAY> </GATEWAY>
    </custom_property>
    <custom_property>
        <PANORAMA_SERVER> </PANORAMA_SERVER>
    </custom_property>
    <custom_property>
        <DNS_SERVER> </DNS_SERVER>
    </custom_property>
    <custom_property>
        <VM_AUTH_KEY> </VM_AUTH_KEY>
    </custom_property>
    <default_profile>VM-50</default_profile>
    <profiles>
        <profile>
                <name>VM-50</name>
                <description>VM-50 profile</description>
                <vcpus>2</vcpus>
                <memory_mb>5120</memory_mb>
                <root_disk_mb>60000</root_disk_mb>
        </profile>
        <profile>
                <name>VM-100-n-200</name>
                <description>VM-100 and VM-200 profile</
description>
                <vcpus>2</vcpus>
                <memory_mb>7168</memory_mb>
                <root_disk_mb>60000</root_disk_mb>
        </profile>
        <profile>
                <name>VM-300</name>
                <description>VM-300 profile</description>
                <vcpus>2</vcpus>
                <memory_mb>9216</memory_mb>
                <root_disk_mb>60000</root_disk_mb>
        </profile>
        <profile>
                <name>VM-1000-HV</name>
                <description>VM-1000-HV profile</description>
                <vcpus>4</vcpus>
                <memory_mb>9216</memory_mb>
                <root_disk_mb>60000</root_disk_mb>
        </profile>
        <profile>
                <name>VM-500</name>
                <description>VM-500 profile</description>
                <vcpus>4</vcpus>
                <memory_mb>16384</memory_mb>
                <root_disk_mb>60000</root_disk_mb>
        </profile>
    </profiles>
    <cdrom>true</cdrom>
```

 ©2024 Palo Alto Networks, Inc.

```
    <bootstrap_file_1>/config/init-cfg.txt</bootstrap_file_1>
    <bootstrap_file_2>/config/bootstrap.xml</bootstrap_file_2>
    <bootstrap_file_3>/license/authcodes</bootstrap_file_3>
</image_properties>
```

STEP 6 | Download the image packaging utility.

1. Log in to the Enterprise NFVIS user interface and select **VM Life Cycle** > **Image Repository**.
2. Click the **Browse Datastore** tab, and navigate to **data** > **intdatastore** > **uploads** > **vmpackagingutility**.
3. Download `nfvisvmpackagingtool.tar` to the conversion folder.
4. Untar the file:

```
tar -xvf nfvisvmpackagingtool.tar
```

STEP 7 | In the conversion folder that contains the qcow2, the `init-config.txt` and the `authcodes` file, run the `nfvpt.py` script. See the `nfvpt.py` image packaging utility documentation.

The following sample creates the image file Palo-Alto-9.1.0, and a VM-100 profile. Options are space-separated (the sample shows options on separate lines for clarity only) and custom options are key-value pairs with a colon separator.

```
./nfvpt.py -o Palo-Alto-9.1.0 -i PA-VM-KVM-9.1.0.qcow2
-n PAN902 -t FIREWALL -r 9.1.0
--monitored false
--privileged true
--bootstrap /config/init-cfg.txt:init-cfg.txt,/license/
authcodes:authcodes
--min_vcpu 2 --max_vcpu 8
--min_mem 4096 --max_mem 16384
--min_disk 10 --max_disk 70
--vnic_max 8
--optimize true
--console_type_serial true
--profile VM-100,"VM-100 profile",2,7168,61440
--default_profile VM-100
--custom HOSTNAME:hello
--custom IP_ADDRESS:10.2.218.24
--custom NETMASK:255.255.255.0
--custom GATEWAY:10.2.218.1
--custom DNS_SERVER:10.55.66.10
--custom PANORAMA_SERVER:0.10.10.0
--custom VM_AUTH_KEY:123451234512345
```

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Upload the converted image.

1. In the NFVIS user interface, select **VM Life Cycle** > **Image Repository** and click the blue Images icon to show the **Drop Files or Click** circle.
2. Drag the converted file into the circle, or click to browse and select your file.
3. In the Status column, click **Start**.

   When the upload is complete, the image is registered, and the file you uploaded displays in the **Image Registration** tab **Images** list.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on Cisco ENCS

Before you begin to deploy the firewall, make sure that you have created network connections for management access to the VM-Series firewall. If you are using Panorama, ensure that Panorama has management connectivity to the firewall.

STEP 1 | Deploy the VM-Series firewall.

1. In Enterprise NFVIS, click **VM Life Cycle** > **Deploy**.
2. Drag the firewall icon to the appropriate network. In this example, the firewall connects to a management network and a LAN network.

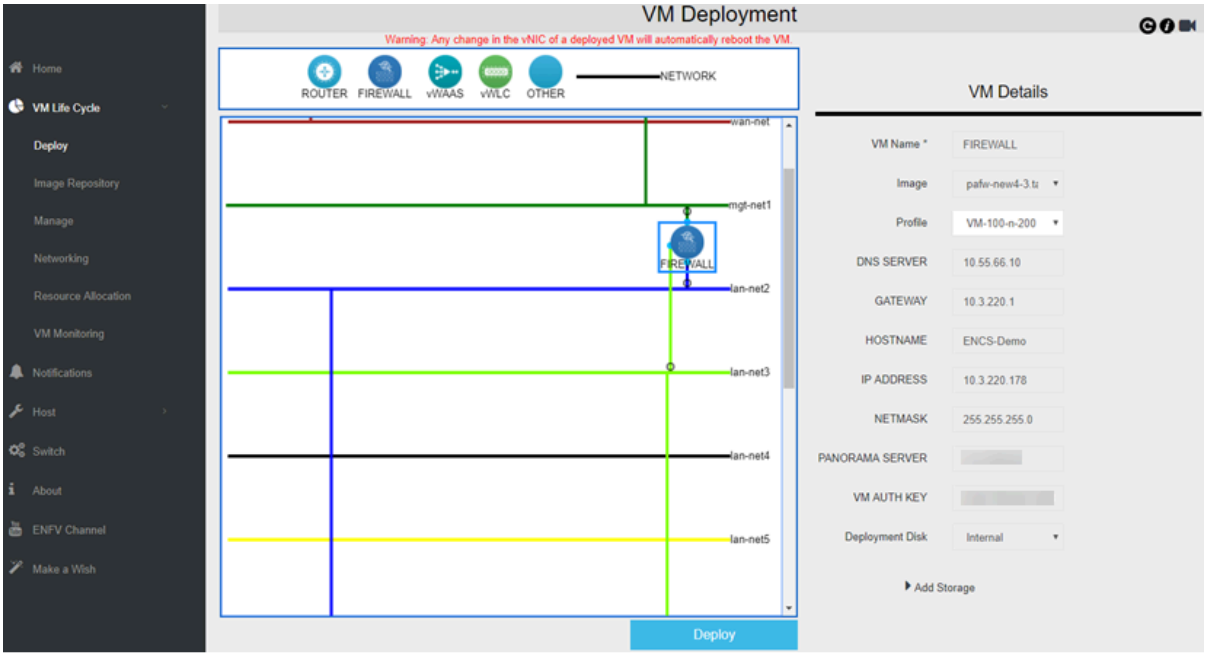

3. **Deploy** the VM-Series firewall.

   If you are using Panorama to manage the firewall, the firewall displays as **Connected** on **Panorama** > **Managed Devices** > **Summary**. If the firewall is not connected to Panorama, check that you have provided the correct Panorama IP address and that the devices can communicate over the network.

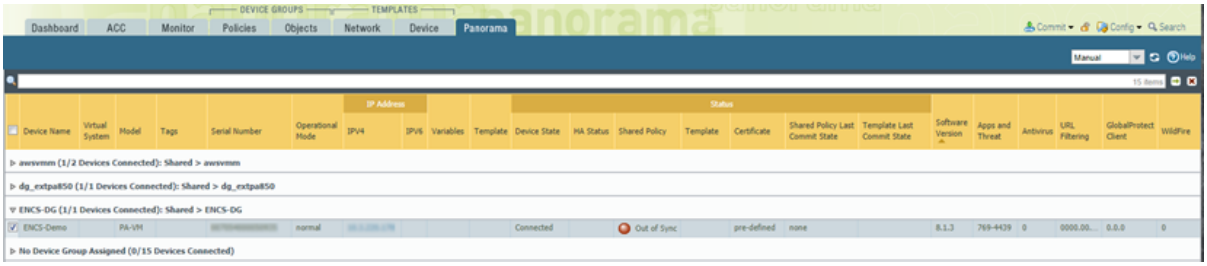

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure the VM-Series firewall dataplane interfaces.

See configure a Layer 3 interface, configure a Layer 2 interface, or configure virtual wires. If using Panorama, the following steps show you how to configure the firewall for a Layer 3 deployment.

1. Add a template and assign the firewall to the template.
2. Select the **Network** and in the Template drop-down, select the template you created.
3. Select **Network** > **Interfaces** > **Ethernet**.
4. Click **ethernet 1/1** and configure as follows:
   - Set **Interface Type** to **Layer3**.
   - On the **Config** tab, assign the interface to the default router.
   - Also on the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone called **UnTrust** for example, and then click **OK**.
   - On the **IPv4** tab, select **DHCP Client** or **Static**. If you choose static, enter the IP address.

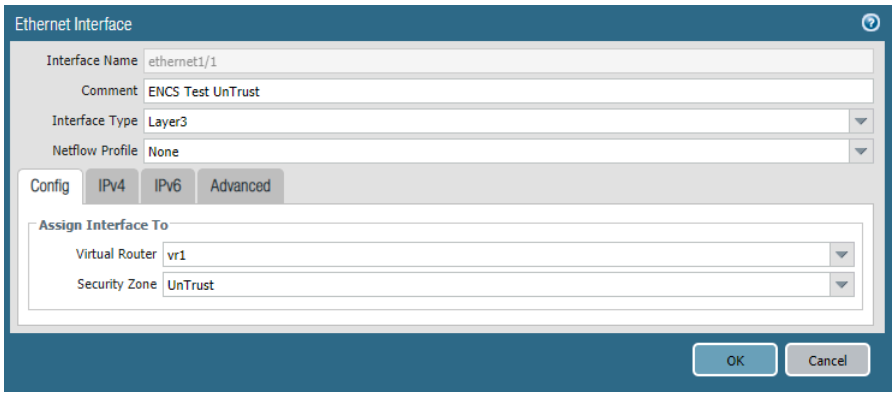

5. Repeat b-e for each network interface.
6. **Commit** > **Commit and Push** to commit all configuration changes to Panorama and the managed firewalls.

   Verify that the link state for the firewall interfaces is up.

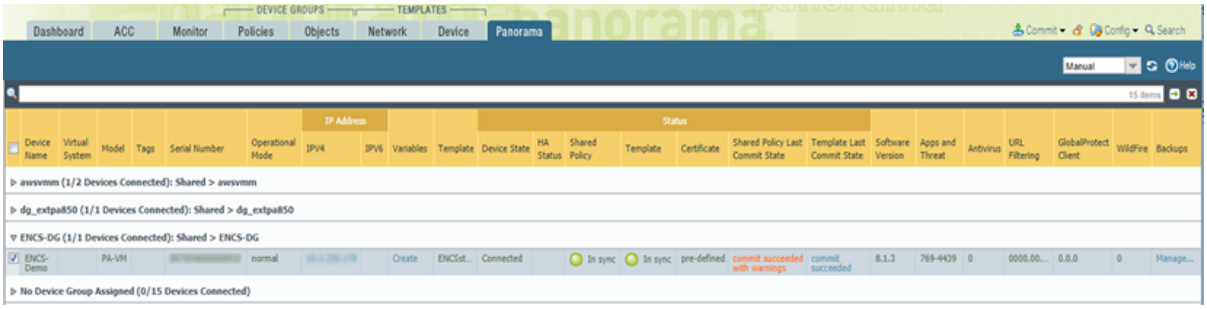

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Configure Security policies to safely enable applications and users on your network.

If using Panorama, the following steps show you how to use device groups to centrally manage policy rules for your managed firewalls.

1. Add a device group and assign the managed firewalls to your device group.

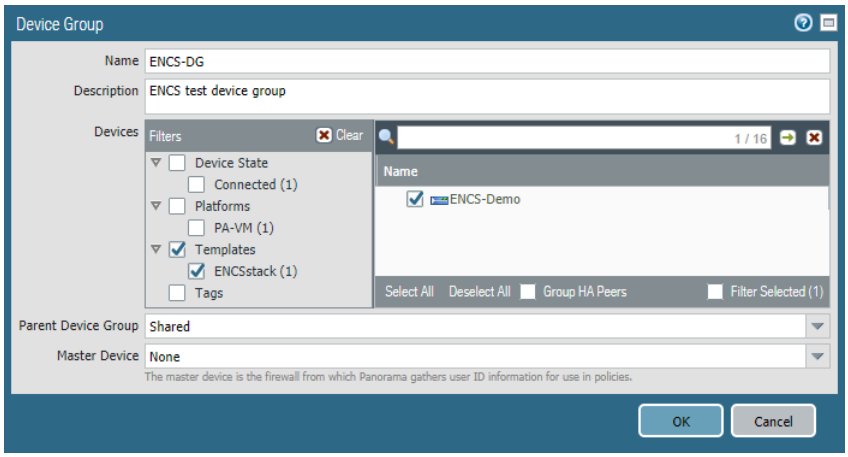

2. Configure the security policies for the device group.

STEP 4 | Verify that the VM-Series firewall is securing traffic on your network.

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set up the VM-Series Firewall on Oracle Cloud Infrastructure

Deploy the VM-Series firewall on Oracle Cloud Infrastructure (OCI) cloud. With the VM-Series on OCI, you can protect and segment your workloads, prevent advanced threats, and improve visibility into your applications as you move to the cloud.

OCI is a public cloud computing service that enables you to run your applications in a highly-available, hosted environment offered by Oracle. You can deploy the VM-Series firewall to secure your applications and services running your OCI environment.

- OCI Shape Types
- Deployments Supported on OCI
- Prepare to Set Up the VM-Series Firewall on OCI
- Deploy the VM-Series Firewall From the Oracle Cloud Marketplace
- Configure Active/Passive HA on OCI

OCI Shape Types

The VM-Series firewalls support the following OCI VM shapes. See Oracle Cloud Infrastructure documentation for more information about VM shapes.

VM-Series Model	Minimum OCI Shape
• VM-100 • Software NFGW Credit-based VM-Series	VM.Standard2.4
• VM-300 • Software NFGW Credit-based VM-Series	VM.Standard2.4
• VM-500 • Software NFGW Credit-based VM-Series	VM.Standard2.8
• VM-700 • Software NFGW Credit-based VM-Series	VM.Standard2.16
• VM-100, VM-300, VM-500, and VM-700 • Software NFGW Credit-based VM-Series	VM.Optimized3.Flex VM.Standard3.Flex

You can deploy the VM-Series firewall on an OCI instance with more resources than the minimum VM-Series System Requirements. If you chooses a larger shape size for the VM-Series firewall model. Although the firewall only uses the maximum vCPUs cores and memory listed on the system requirements page, it does take advantage of the faster network performance that the larger shape provides.

 ©2024 Palo Alto Networks, Inc.

Deployments Supported on OCI

Use the VM-Series firewall on OCI to secure your cloud environment in the following scenarios:

- North-South Traffic—You can use the VM-Series firewall to secure traffic entering your cloud network from an untrusted source or exiting your cloud network to reach an untrusted source. For either type of traffic, you must configure route table rules in your Virtual Cloud Network (VCN) and NAT policy rules on the firewall.

  In this example, outbound traffic is exiting the trust subnet in your VCN. You must configure source address translation policy onto a public IP address and a route table rule that redirects that traffic to the firewall. The route rule points outgoing traffic to the firewall's interface in the trust subnet of the VCN. When the firewall receives this traffic, it performs the source address translation on the traffic and applies any other security policy you have configured.

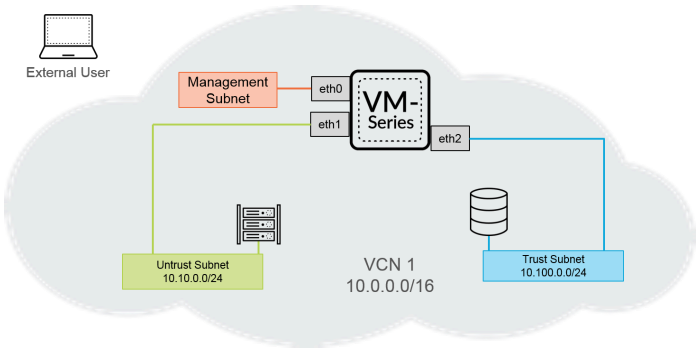

- Inter-VCN Traffic (East-West)—The VM-Series firewall allows you to secure traffic moving within your cloud environment between VCNs. Each subnet must belong to a different VCN because, by default, no route rules are used to enable traffic within a VCN. In this scenario, you configure an interface on the firewall connected to a subnet in each VCN.

  In the example below, a user in the Trust Subnet wants to access data in the DB Subnet. Configure a route on OCI that reaches DB Subnet CIDR next hop, which points to the interface Trust Subnet network on the VM-Series firewall.

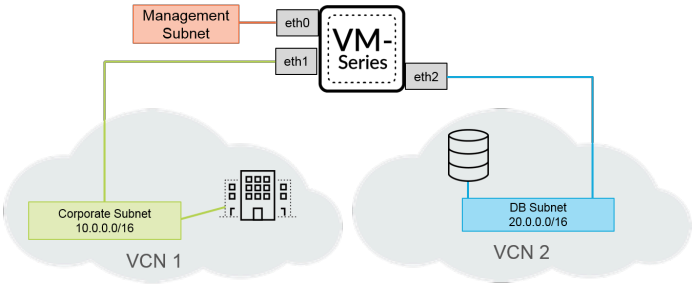

 ©2024 Palo Alto Networks, Inc.

Prepare to Set Up the VM-Series Firewall on OCI

The process to deploy the VM-Series firewall on Oracle Cloud Infrastructure requires the completion of preparation tasks.

- Virtual Cloud Networks
- SSH Keys
- Initial Configuration User Data

Virtual Cloud Networks

A virtual cloud network (VCN) is a virtual, private network that you set up in your OCI environment. To deploy the VM-Series firewall in OCI, your VCN must have at least three virtual network interfaces cards (VNICs) for the management interface and two data interfaces.

OCI uses a series of route tables to send traffic out of your VCN and one route table is added to each subnet. A subnet is a division of your VCN. If you do not specify a route table, the subnet uses the VCN's default route table.Each route table rule specifies a destination CIDR block and a next hop (target) for any traffic that matches the CIDR. OCI only uses a subnet's route table if the destination IP address is outside the VCN's specified CIDR block; route rules are not required to enable traffic within the VCN. And, if traffic has overlapping rules, OCI use the most specific rule in the route table to route traffic.

If there is no route rule that matches the traffic that is attempting to leave the VCN, the traffic is dropped.

Each subnet requires a route table and once you have added a route table to a subnet, you cannot change it. However, you can add, remove, or edit rules in a route table after it has been created.

SSH Keys

You must create an SSH key pair to login to the firewall for the first time. You cannot use the default username and password to access the firewall for the first time. After the firewall boots up for the first time, you must access the firewall through the CLI and create a new username and password.

1. Create an SSH key pair and store the SSH Key pair in the default location for your operating system.
   - On Linux or MacOS, use `ssh-keygen` to create the key pair in your .ssh directory.
   - On Windows, use PuTTYgen to create the key pair.

     The content of the **Key comment** field does not matter to the VM-Series firewall; you can accept the default (the key creation date) or enter a comment that helps you remember the name of the key pair. Use the **Save private key** button to store the private key in your .ssh directory.

 ©2024 Palo Alto Networks, Inc.

2. Select the full public key.
   - Linux or MacOS:

     Open your public key in a text editor and copy the public key.
   - Windows: You must use the PuTTY Key Generator to view the public key. Launch PuTTYgen, click Load, and browse to private key you saved in your .ssh directory.

     In PuTTYgen, scroll down to ensure you select the entire key, right click, and choose Copy.

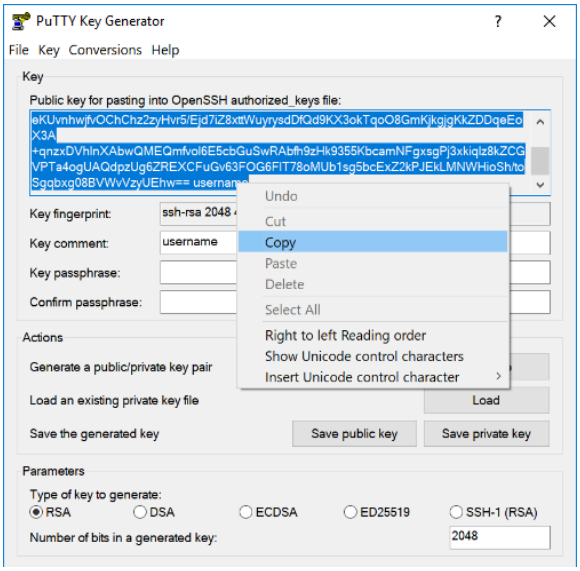

Initial Configuration User Data

You must provide the following bootstrapping parameters when setting up the VM-Series firewall instance. OCI uses this information to perform the initial configuration of the firewall, which provides the firewall with a hostname and license and connects the firewall to Panorama, if applicable.

The Panorama-related fields are required only if you have a Panorama appliance and want use Panorama to manage your VM-Series firewall.

Field	Description
hostname=	Host name for the firewall.
vm-auth-key=	Virtual machine authentication key for registering the firewall with Panorama.
panorama-server=	IPv4 or IPv6 address of the primary Panorama server. This field is not required but recommended for centrally managing your firewalls.
panorama-server-2=	IPv4 or IPv6 address of the secondary Panorama server. This field is not required but recommended.

 ©2024 Palo Alto Networks, Inc.

Field	Description
tplname=	Panorama template stack name. If you add a Panorama server IP address, as a best practice assign the firewall to a template stack on Panorama and enter the template stack name in this field so that you can centrally manage and push configuration settings to the firewall.
dgname=	Panorama device group name. If you add a Panorama server IP address, as a best practice create a device group on Panorama and enter the device group name in this field so that you can group the firewalls logically and push policy rules to the firewall.
authcodes=	Used to license the VM-Series firewall with the Palo Alto Networks licensing server.
op-command-modes=jumbo-frame	Used to enable jumbo frame mode on the VM-Series firewall. Because OCI deploys VM instances in jumbo mode by default, it is recommended that you launch the VM-Series firewall in jumbo mode to achieve the best throughput.

Paste the bootstrapping parameters into the OCI console in the following format.

hostname=<*fw-hostname*>

vm-auth-key=<*auth-key*>

panorama-server=<*panorama-ip*>

panorama-server-2=<*panorama2-ip*>

tplname=<*template-stack-name*>

dgname=<*device-group-name*>

authocodes=<*firewall-authcode*>

op-command-modes=jumbo-frame

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall From the Oracle Cloud Marketplace

Complete the following procedure to deploy the VM-Series firewall in OCI from the Oracle Cloud Marketplace.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Log in to the Oracle Cloud Marketplace.

STEP 2 | Find the VM-Series firewall application in the Oracle Cloud Marketplace.

1. Search for Palo Alto Networks and a list of offerings for the VM-Series firewall will display.
2. Select an offering.
3. Click **Get App**.
4. Select your **Region** and click **Sign In**.
5. Select the **Version** and **Compartment**.
6. Accept the Oracle and Partner terms.
7. Click **Launch Instance**.

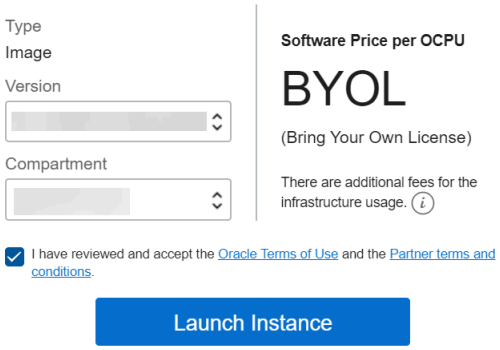

STEP 3 | Enter a descriptive **Name** for your VM-Series firewall instance.

STEP 4 | Select an **Availability Domain**.

STEP 5 | Select **Virtual Machine** under **Shape Type**.

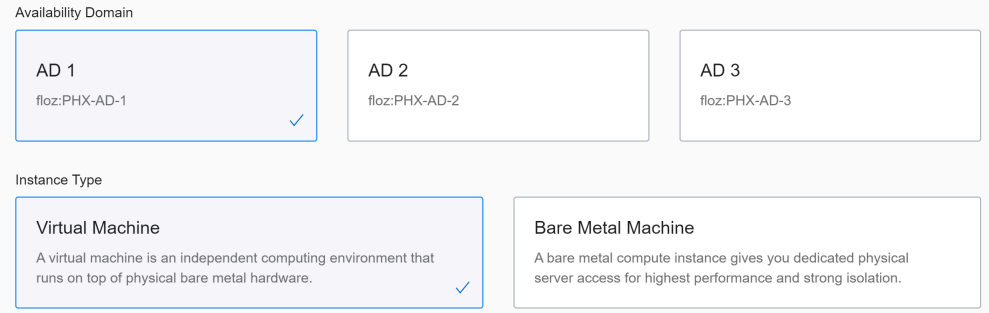

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Select the shape with the number of CPUs, amount of RAM, and number of interfaces required for the VM-Series firewall model. See the Compute Shapes page for the amount resources provided by the different compute shapes. See VM-Series System Requirements for more information about the resources required for each VM-Series firewall model.

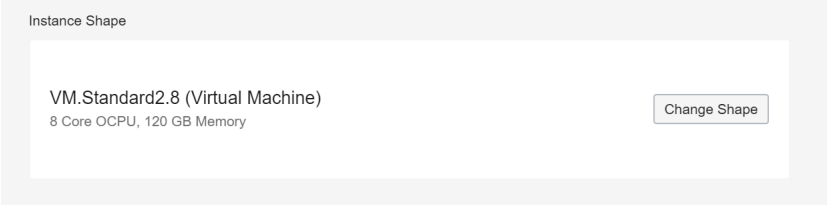

STEP 7 | Under Networking, select your **Virtual cloud network compartment**, **Virtual cloud network**, **Subnet compartment**, and **Subnet** for your management interface. You can only add one interface when creating the VM-Series firewall instance. You will add additional interfaces later.

STEP 8 | (Optional) Set the boot volume to a size larger than the default. By default, the boot volume is set to 60GB. Complete this procedure if you require a larger boot volume to support features such as attaching logs.

1. Select **Custom boot volume size (in GB)**.
2. Enter 60 or greater. 60 GB is the minimum hard drive size required by the VM-Series firewall.

STEP 9 | Add your SSH key.

1. Under **Add SSH Key**, select **Paste SSH Key**.
2. Paste your SSH key into the field provided.

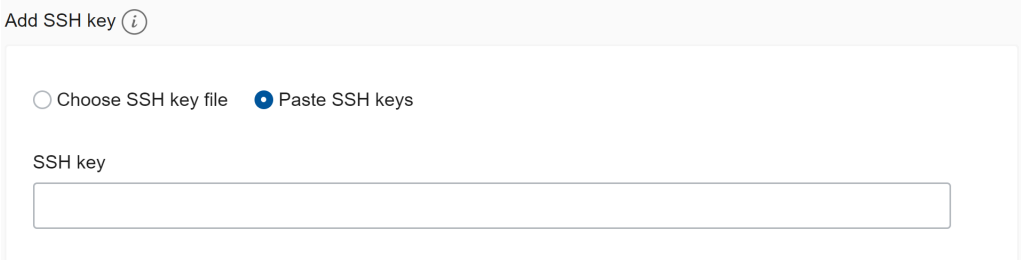

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Add the bootstrapping parameters.

1. Click **Show Advanced Options**.
2. Under **User data**, select **Paste cloud-init script**.
3. Paste the boostrap parameters into the field provided.

   hostname=<*fw-hostname*>

   vm-auth-key=<*auth-key*>

   panorama-server=<*panorama-ip*>

   panorama-server-2=<*panorama2-ip*>

   tplname=<*template-stack-name*>

   dgname=<*device-group-name*>

   authcodes=<*firewall-authcode*>

   op-command-modes=jumbo-frame

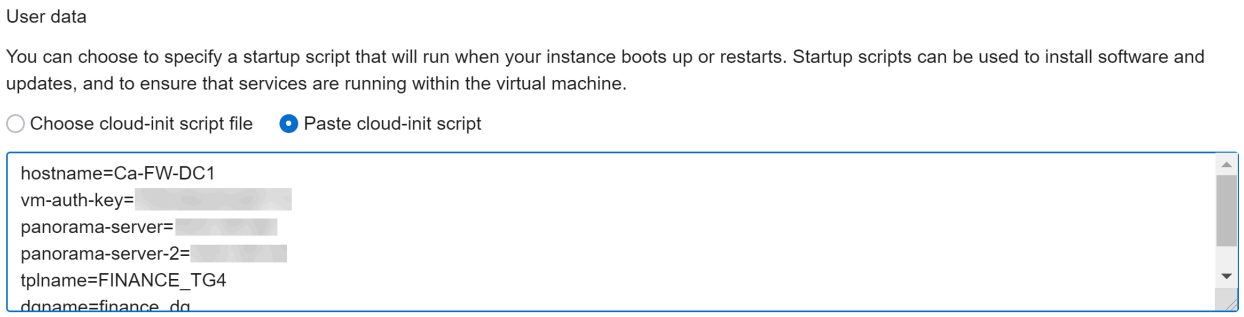

STEP 11 | Click **Create**.

When the VM-Series firewall is launched, OCI creates and attaches a primary VNIC to the instance. This VNIC resides in the subnet you specified in the instance network setting and connects to the VM-Series firewall's management interface.

STEP 12 | Configure a new administrative password for the firewall.

1. Use the management IP address to SSH into the command line interface (CLI) of the VM-Series firewall.
2. Enter the following command to log in to the firewall:

   ssh-i <*private_key.pem*> admin@<*public-ip_address*>
3. Configure a new password, using the following command and follow the onscreen prompts:

   configure

   set mgt-config users admin password

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Attach a vNIC to your VM-Series firewall instance for each data interface. You must attach at least two data interfaces to your firewall instance—untrust and trust.

1. Select your newly launched VM-Series firewall instance and select **Attached VNICs** > **Create VNIC**.
2. Enter a descriptive **Name** for your vNIC.
3. Select your VCN from the **Virtual Cloud Network** drop-down.
4. Select your subnet from the **Subnet** drop-down.
5. Specify a **Private IP Address**. This is only required if your want to choose a particular IP for the vNIC. If you do not specify an IP, OCI will assign an IP address from the CIDR block you assigned to the subnet.
6. Select **Assign Public IP Address** for public facing vNICs such as your untrust subnet.
7. Click **Create VNIC**.
8. Repeat this procedure for each vNIC your deployment requires.

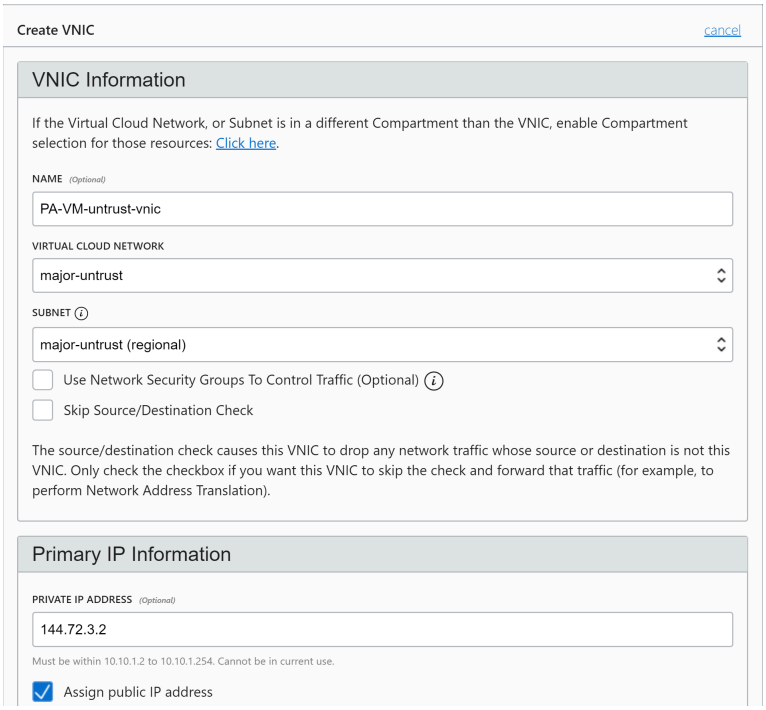

 ©2024 Palo Alto Networks, Inc.

STEP 14 | Configure the dataplane network interfaces as Layer 3 interfaces on the firewall.

1. Log in to the firewall.
2. Select **Network** > **Interfaces** > **Ethernet**.
3. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example untrust-zone, and then click **OK**.
   - On the **IPv4** tab, select either **Static**.
   - Click **Add** in the IP section and enter the IP address and network mask for the interface. Make sure that the IP address matches the IP address that you assigned to the corresponding subnet in VCN. For example, if you add this interface to your untrust zone, make sure you assign the untrust vNIC IP address configured in your VCN.
4. Repeat this procedure for each vNIC configured in your VCN except your management vNIC.

Always only delete interfaces at the bottom of the interface list. Deleting firewall interfaces in the wrong order results in a interface mismatch between the firewall and OCI. For example, say you have five data interfaces, then delete interface two on the firewall and add a new interface at the bottom. After rebooting the firewall, the newly added interface will take the place of the deleted interface two instead of taking a place at the bottom of the list.

 ©2024 Palo Alto Networks, Inc.

Configure Active/Passive HA on OCI

You can configure a pair of VM-Series firewalls on OCI in an active/passive high availability (HA) configuration. To ensure uptime in an HA setup on OCI, you must create a secondary, floating IP addresses that can quickly move from one peer to the other. When the active firewall goes down, the floating IP address moves from the active to the passive firewall so that the passive firewall can seamlessly secure traffic as soon as it becomes the active peer. In addition to the floating IP address, the HA peers also need HA links—a control link (HA1) and a data link (HA2)—to synchronize data and maintain state information.

The VM-Series firewall for OCI in FIPS mode does not support high availability.

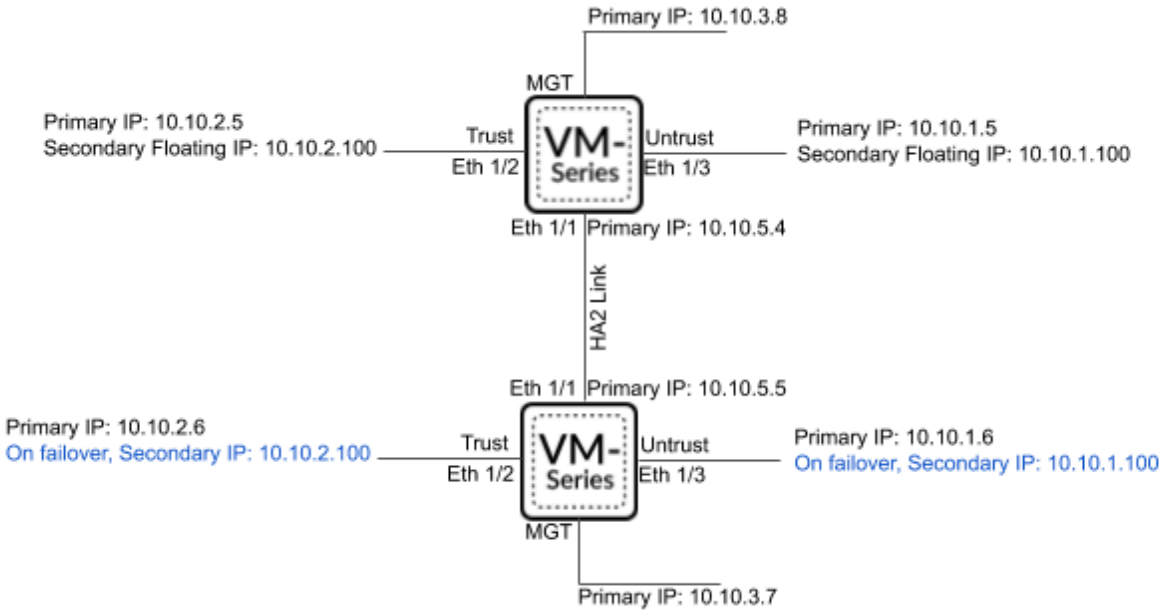

To allow the firewalls to move the floating IP address upon failover, you must place the firewall instances in a dynamic group on OCI. Dynamic groups allow you to group the firewall instances as principal actors and create policy to allow the instances in the dynamic group to make API calls against OCI services. You will use matching rules to add the HA peer instances to the dynamic group and then create the policy the floating IP from one VNIC to another.

Both VM-Series firewalls in the HA pair must have the same number of network interfaces. Each firewall requires a minimum of four interfaces—management, untrust, trust, and HA. You can configure additional data interfaces as required by your deployment.

- **Management interface**—the private and public IP addresses associated with the primary interface. You can use the private IP address on the management interface as the IP address for the HA1 interface between the peers. If you want a dedicated HA interface, you must attach an additional interface to each firewall, for a total of five interfaces each.
- **Untrust** and **trust interfaces**—each of these data interfaces on the active HA peer require a primary and secondary IP address. Upon failover, when the passive HA peer transitions to the

 ©2024 Palo Alto Networks, Inc.

active state, the secondary private IP address is detached from the previously active peer and attached to the now active HA peer.

- **HA2 interface**—this interface has a single private IP address on each HA peer. The HA2 interface is the data link peers use to synchronize sessions, forwarding tables, IPsec security associations, and ARP tables.

STEP 1 | Deploy the VM-Series Firewall From the Oracle Cloud Marketplace and set up the network interfaces for HA.

1. (Optional) Configure a dedicated HA1 interface on each HA peer.
    1. From the OCI Console, select **Compute** > **Instances** and click on the name of your active peer instance.
    2. Select **Attached VNICs** and click **Create VNIC**.
    3. Enter a descriptive name for your HA1 interface.
    4. Select the VCN and subnet.
    5. Enter a private IP address.
    6. Click **Create VNIC**.
    7. Repeat this process on your passive peer instance.
2. Configure an HA2 interface on each HA peer.
    1. From the OCI Console, select **Compute** > **Instances** and click on the name of your active peer instance.
    2. Select **Attached VNICs** and click **Create VNIC**.
    3. Enter a descriptive name for your HA2 interface.
    4. Select the VCN and subnet. The HA2 interface should be on a separate subnet from your data interfaces.
    5. Enter a private IP address.
    6. Click **Create VNIC**.
    7. Repeat this process on your passive peer instance.
3. Add a secondary IP address to your dataplane interfaces on the active peer.
    1. From the OCI Console, select **Compute** > **Instances** and click on the name of your active peer instance.
    2. Select **Attached VNICs** and click on your untrust VNIC.
    3. Select **IP Addresses** and click **Assign Private IP Address**.
    4. Enter the IP address and click **Assign**.
    5. Repeat this procedure for each dataplane interface on your active peer.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Create security rules to allow the HA peers to synchronize data and maintain state information. By default, OCI allows ICMP traffic only. You must open the necessary HA ports.

1. Open the ports for your HA1 interface.
   1. From the OCI Console, select **Networking** > **Virtual Cloud Networks** and select your VCN.
   2. Select **Subnets** and select the subnet containing your HA1 interface.
   3. Select **Security Lists** and click the default security list to edit it.
   4. Click **Add Ingress Rule**.
   5. Enter the **Source CIDR** that includes the HA peer HA1 port IP address.
   6. Select **TCP** from the **IP Protocol** drop-down.
   7. Click **+Additional Ingress Rule**. You need to create two additional rules for TCP ports 28260 and 28769.
   8. If encryption is enabled on your VM-Series firewall for the HA1 link, create an additional rules for ICMP and TCP port 28.
   9. Click **Add Ingress Rules**.

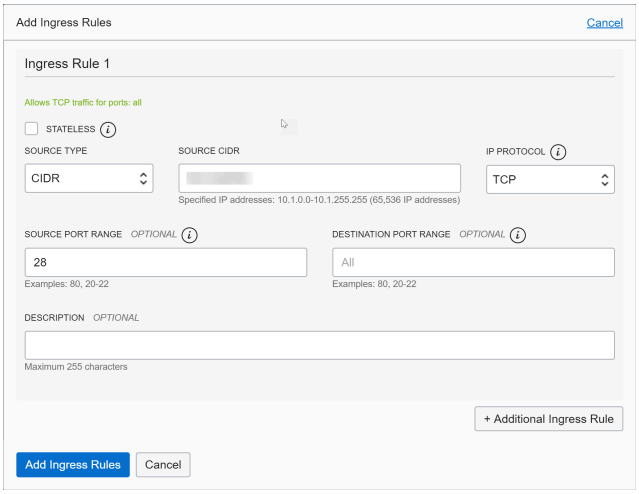

2. Open the ports for your HA2 interface.
   1. From the OCI Console, select **Networking** > **Virtual Cloud Networks** and select your VCN.
   2. Select **Subnets** and select the subnet containing your HA2 interface.
   3. Select **Security Lists** and click the default security list to edit it.
   4. Click **Add Ingress Rule**.
   5. Enter the **Source CIDR** that includes the HA peer HA2 port IP address.
   6. Select **UDP** or **IP** from the **IP Protocol** drop-down.
   7. If the transport mode is UDP, enter **29281** into **Source Port Name**. If the transport mode is IP, enter **99** into **Source Port Name**.
   8. Click **Add Ingress Rules**.

 ©2024 Palo Alto Networks, Inc.

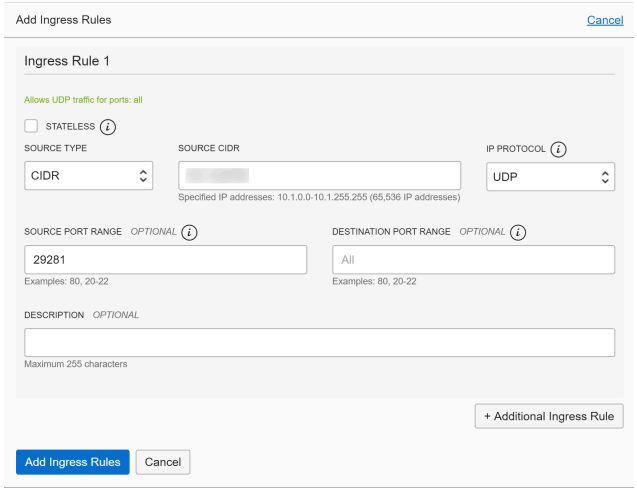

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Add both HA peers to a dynamic group and create policy that allows the HA peers to move the floating IP address. You must have the OCID of each HA peer instance to build the dynamic group matching rules, so have those on hand to past into the rule builder.

1. Create the dynamic group.
   1. From the OCI Console, select **Identity** > **Dynamic Groups** > **Create Dynamic Group**.
   2. Enter a descriptive **Name** for your dynamic group.
   3. Click **Rule Builder**.
   4. Select **Any of the following rules** from the first drop-down.
   5. Select **Match instances with ID:** from the **Attributes** drop-down and paste one of the peer OCIDs into the **Value** field.
   6. Click **+Additional Line**.
   7. Select **Match instances with ID:** from the **Attributes** drop-down and paste the other peer OCID into the **Value** field.
   8. Click **Add Rule**.

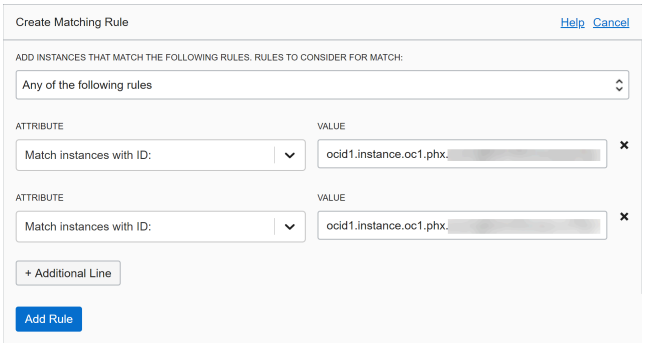

   9. Click **Create Dynamic Group**.
2. Create the policy rule.
   1. From the OCI Console, select **Identity** > **Policies** > **Create Policy**.
   2. Enter a descriptive **Name** for your policy.
   3. Enter the first policy statement.

      Allow dynamic-group <*dynamic_group_name*> to use virtual-network-family in compartment <*compartment_name*>
   4. Click **+Another Statement**.
   5. Enter the second policy statement.

      Allow dynamic-group <*dynamic_group_name*> to use instance-family in compartment <*compartment_name*>
   6. Click **Create**.

 ©2024 Palo Alto Networks, Inc.

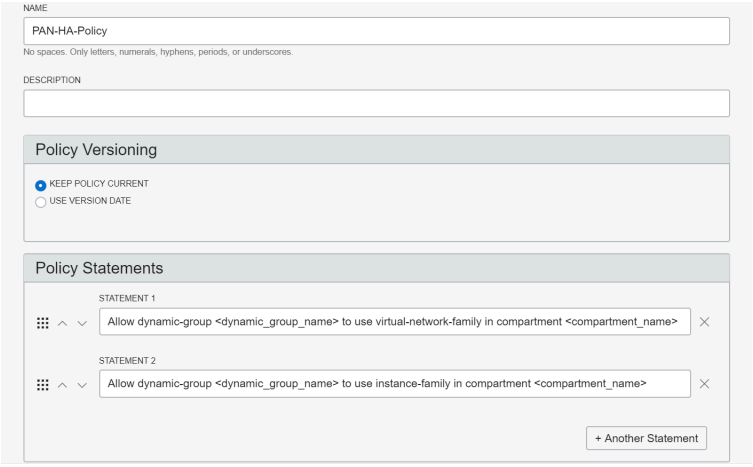

STEP 4 | Configure the interfaces on the firewall. You must configure the HA2 data link and at least two Layer 3 interfaces for your untrust and trust interfaces. Complete this workflow on the first HA peer and then repeat the steps on the second HA peer.

1. Log in to the firewall web interface.
2. (Optional) If you are using the management interface as HA1, you must set the interface IP Type to static and configure a DNS server.
   1. Select **Device** > **Setup** > **Interfaces** > **Management**.
   2. Set the **IP Type** to **Static**.
   3. Enter the private **IP address** of the primary VNIC of your VM-Series firewall instance.
   4. Click **OK**.
   5. Select **Device** > **Setup** > **Services**.
   6. Click **Edit**.
   7. Enter the IP address of the **Primary DNS Server**.
   8. Click **OK**.
   9. **Commit** your changes.
3. Select **Network** > **Interfaces** > **Ethernet** and click on your untrust interface. In this example, the HA2 interface is 1/1, the trust interface is ethernet 1/2, and the untrust interface is ethernet 1/3.
4. Click the link for **ethernet 1/1** and configure as follows:
   - **Interface Type**: **HA**
5. Click the link for **ethernet 1/2** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example trust-zone, and then click **OK**.
   - On the **IPv4** tab, select either **Static**.
   - Click **Add** in the IP section and enter the primary IP address and network mask for the interface. Make sure that the IP address matches the IP address that you assigned to

 ©2024 Palo Alto Networks, Inc.

the corresponding subnet in VCN. For example, if you add this interface to your trust zone, make sure you assign the trust vNIC IP address configured in your VCN.

- Click **Add** in the IP section and enter the secondary, floating IP address and network mask.

6. Click the link for **ethernet 1/3** and configure as follows:
   - **Interface Type**: **Layer3**
   - On the **Config** tab, assign the interface to the default router.
   - On the **Config** tab, expand the **Security Zone** drop-down and select **New Zone**. Define a new zone, for example untrust-zone, and then click **OK**.
   - On the **IPv4** tab, select either **Static**.
   - Click **Add** in the IP section and enter the primary IP address and network mask for the interface. Make sure that the IP address matches the IP address that you assigned to the corresponding subnet in VCN. For example, if you add this interface to your untrust zone, make sure you assign the untrust vNIC IP address configured in your VCN.
   - Click **Add** in the IP section and enter the secondary, floating IP address and network mask.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Enable HA.

1. Select **Device** > **High Availability** > **General**.
2. Edit the Setup settings.
3. Enter the private IP address of the passive peer in the **Peer HA1 IP address field**.
4. Click **OK**.

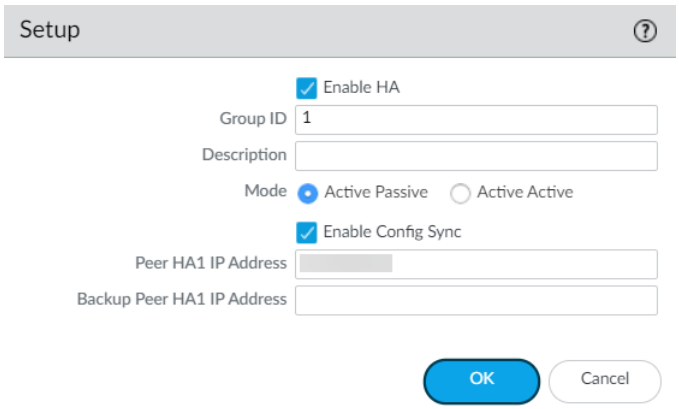

5. (Optional) Edit the Control Link (HA1). If you do not plan to use the management interface for the control link and have added an additional interface (for example ethernet 1/4), edit this section to select the interface to use for HA1 communication.
6. Edit the Data Link (HA2) to use **Port** ethernet 1/1 and add the IP address of active peer and the **Gateway** IP address for the subnet.
7. Select **IP** or **UDP** from the Transport drop-down. Ethernet is not supported.

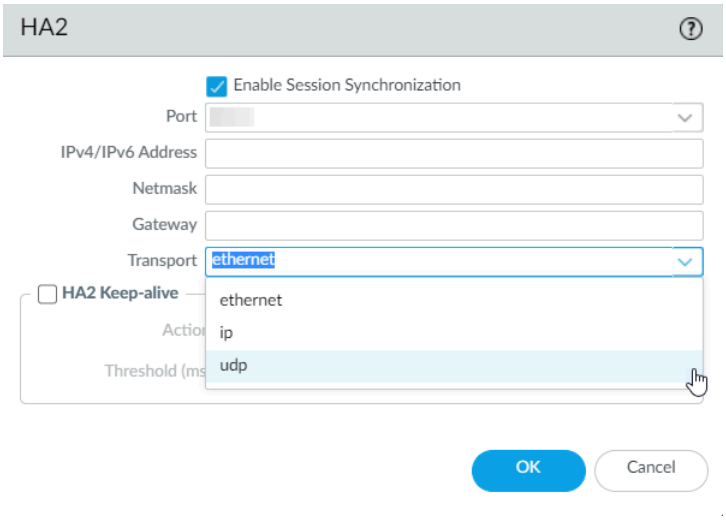

8. Click **OK**.

STEP 6 | **Commit** your changes.

STEP 7 | Repeat step 4 and step 5 on the passive HA peer.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | After your finish configuring HA on both firewalls, verify that the firewalls are paired in active/passive HA.

1. Access the **Dashboard** on both firewalls and view the High Availability widget.
2. On the active HA peer, click **Sync to peer**.
3. Confirm that the firewalls are paired and synced.
   - On the passive firewall: the state of the local firewall should display **Passive** and the **Running Config** should show as Synchronized.
   - On the active firewall: the state of the local firewall should display **Active** and the **Running Config** should show as Synchronized.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on Alibaba Cloud

Deploying the VM-Series firewall on Alibaba Cloud protects networks you create within Alibaba Cloud. You can deploy VM-Series firewalls to protect internet facing applications and hybrid cloud deployments.

- VM-Series Firewall on Alibaba Cloud
- Minimum System Requirements for the VM-Series Firewall on Alibaba Cloud
- Prepare to Deploy the VM-Series Firewall on Alibaba Cloud
- Deploy the VM-Series Firewall on Alibaba Cloud

VM-Series Firewall on Alibaba Cloud

You can deploy the VM-Series firewall to secure inbound and outbound north-south traffic in Alibaba Cloud.

Securing east-west traffic within the same VPC is not supported because Alibaba Cloud does not support subnet routing.

The VM-Series firewall on Alibaba Cloud runs on the KVM hypervisor and supports up to 8 network interfaces when you select an Alibaba Cloud instance with sufficient resources (see Minimum System Requirements for the VM-Series Firewall on Alibaba Cloud).

The VM-Series firewall on Alibaba Cloud supports BYOL licensing and the VM-Series ELA on Alibaba Cloud International Regions and Mainland China. PAYG licensing is not currently supported.

In Alibaba Cloud, your VPC logically isolates your virtual network. After creating a VPC, you can create VSwitches to further segment your virtual private network, as shown in the following diagram. To secure inbound traffic, both DNAT and SNAT must be configured on the firewall.

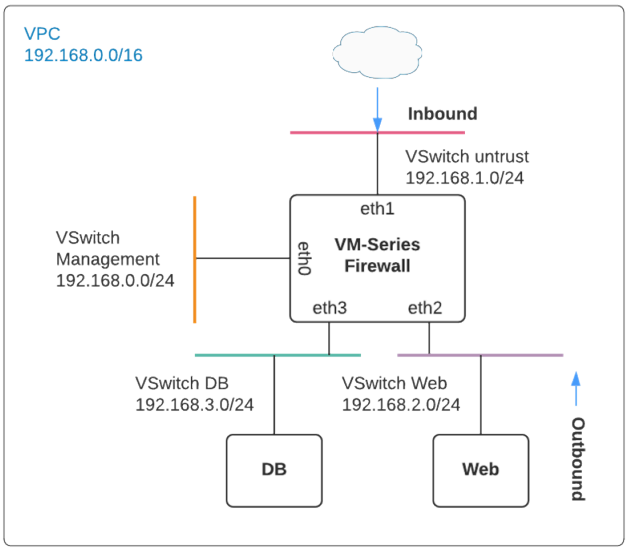

Inbound traffic originates from a client outside of your VPC going to the VM-Series firewall untrust interface. The firewall inspects the traffic and sends it to an application through the trust interface. Traffic returning from the application must travel through the VM-Series firewall trust interface, which inspects the return traffic flow and sends it out through the untrust interface.

Outbound traffic typically originates from an external application. Typically you route the internet facing traffic within a VPC to a NAT gateway (with EIP attached). To do this, add a default gateway route in the VPC routing table, with the VM-Series firewall IP address of the application subnet as the next hop. Configure SNAT using the untrust interface IP to ensure traffic originating from the internet returns through the VM-Series firewall.

Refer to Secure North-South Traffic on Alibaba Cloud for a sample configuration.

 ©2024 Palo Alto Networks, Inc.

Minimum System Requirements for the VM-Series Firewall on Alibaba Cloud

On Alibaba Cloud, you can deploy the VM-Series firewall on the KVM hypervisor (see VM-Series Deployments).

- VM-Series Firewall Software Requirements
- Alibaba Cloud Instance Type Recommendations for the VM-Series Firewall
- Alibaba Cloud CLI

VM-Series Firewall Software Requirements

Ensure that you have the software and licenses required to complete a VM-Series deployment on Alibaba Cloud.

- To deploy the VM-Series firewall on Alibaba Cloud, you must use a VM-Series image you obtain from the Alibaba Marketplace. The image incorporates PAN-OS version 10.0.3 and VM-Series plugin version 2.0.3.
- Before you deploy, choose the VM-Series ELA or BYOL license, a capacity license, and a subscription bundle. See VM-Series Model License Types.
- You must be able to SSH into the VM-Series firewall to complete the deployment. If your OS does not support SSH, install third-party software, such as Putty.

Alibaba Cloud Instance Type Recommendations for the VM-Series Firewall

Before creating the VM-Series firewall, you must choose an Elastic Compute Service (ECS) instance type that supports the minimum system requirements for your VM-Series model. Review the instance type documentation to ensure the ECS instance type has the resources to secure your network configuration.

VM-Series Model	Elastic Compute Service Instance Types
VM-100	ecs.g5.xlarge, ecs.sn2ne.xlarge
VM-300	ecs.g5.xlarge, ecs.sn2ne.xlarge
VM-500	ecs.g5.2xlarge, ecs.sn2ne.2xlarge
VM-700	ecs.g5.4xlarge, ecs.sn2ne.4xlarge

Alibaba Cloud CLI

Aliyun version 3.0.4 or higher. See Prepare to Use the Aliyun Command Line Interface.

 ©2024 Palo Alto Networks, Inc.

Prepare to Deploy the VM-Series Firewall on Alibaba Cloud

This task uses the Aliyun CLI to create a VPC and VSwitches for the VM-Series firewall, however, you should plan your network before you start. Evaluate the applications you want to protect, and determine where you will deploy the VM-Series firewall to inspect and secure north-south traffic.

- Choose Licenses and Plan Networks
- Prepare to Use the Aliyun Command Line Interface

Choose Licenses and Plan Networks

Evaluate the applications you need to protect and create networks that permit the VM-Series firewall to inspect your inbound and outbound application traffic.

STEP 1 | Plan and design your VPC.

1. Plan networks, including CIDR Blocks for your VPCs and VSwitches.

   Refer to Create a VPC and Configure Networks for a sample procedure.
2. Plan your IP addresses. If you need specific addresses or address ranges, refer to the Elastic IP Address User Guide.
3. Plan security groups.

STEP 2 | Evaluate your applications and network configurations and calculate the firewall capacity you need to secure your applications and networks.

STEP 3 | Obtain VM-Series firewall licenses.

Although you do not need a license to install the VM-Series firewall (you can activate a license after the installation), you must choose an appropriate VM-Series model and ECS instance type before deploying the firewall.

1. Choose a VM-Series model.

   The VM-Series firewall supports up to 8 interfaces, provided the VM-Series model and Alibaba Cloud instance have sufficient resources.You can use the model

   Use the VM-Series model you have chosen to choose one of the Alibaba Cloud Instance Type Recommendations for the VM-Series Firewall.
2. Choose a VM-Series capacity license that meets your needs.
3. Purchase a BYOL subscription bundle (if you do not already have one). You receive an auth code for your VM-Series subscription, and you must supply it during the deployment.

STEP 4 | Plan how to configure Alibaba accounts and permissions to access the VM-Series firewall. For a start, see the Security FAQ, and learn about Instance RAM Roles.

 ©2024 Palo Alto Networks, Inc.

Prepare to Use the Aliyun Command Line Interface

This chapter focuses on the ECS Console, however, everything you do in the ECS Console can be done from the Aliyun command line interface. The CLI is required if you want to use the VM-Series firewall to secure load balancing on Alibaba Cloud.

Install and configure a recent version of Aliyun, the Alibaba Cloud command line interface.

STEP 1 | Create an AccessKey and save the Access Key ID and Secret in a secure place.

STEP 2 | Download a supported version of Aliyun from https://github.com/aliyun/aliyun-cli.

STEP 3 | Install Aliyun.

STEP 4 | Configure Aliyun.

The configuration prompts you for your Access Key information and other information.

If your deployment uses a storage bucket, the region must match the region for your bucket.

```
aliyun configure
Configuring profile '' in '' authenticate mode...
Access Key Id [*************8rq]: *************8rq
Access Key Secret [***************************tM2]:
***************************tM2
Default Region Id [us-west-1]: us-west-1
Default Output Format [json]: json (Only support json))
Default Language [zh|en] en: en
Saving profile[] ...Done.
 available regions:
...
```

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on Alibaba Cloud

The VM-Series firewall assumes a minimum of three interfaces: management, untrust, and trust. When you create an Alibaba Cloud VPC, it is logically isolated. To segment your virtual private network into subnets you create VSwitches, each having its own CIDR block. Because the VM-Series firewall has multiple interfaces, it can inspect traffic on all subnets.

Typically external inbound traffic encounters the VM-Series firewall untrust interface. The firewall inspects the inbound traffic and sends it to an application through the trust interface. Return traffic from the application goes to the firewall's trust interface, where the firewall inspects the return traffic and sends it out through the untrust interface.

The following tasks demonstrate how to use the console to create the VM-Series firewall infrastructure.

- Create a VPC and Configure Networks
- Create and Configure the VM-Series Firewall
- Secure North-South Traffic on Alibaba Cloud
- Configure Load Balancing on Alibaba Cloud

Create a VPC and Configure Networks

Use the Alibaba Cloud console to create a VPC, VSwitches, security groups, and security group rules.

All VM-Series firewall interfaces must be assigned an IPv4 address when deployed in a public cloud environment. IPv6 addresses are not supported.

STEP 1 | Open the VPC console and select your region from the menu. Note, the region you select must provide one of the instance types that Palo Alto Networks supports.

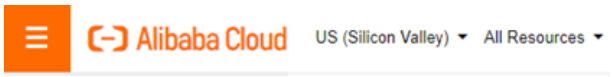

STEP 2 | From the Alibaba Cloud Console home page, select **Products and Services** > **Networking** > **Virtual Private Cloud**.

STEP 3 | **Create a VPC**.

In this step you create a VPC and Management, Untrust, and Trust VSwitches. The ECS console creates a VPC and a switch using the same form.

1. Select **Create VPC**.

   Specify the VPC name, an IPv4 CIDR Block, and a description. Refer to Create a VPC.

Property	Value
Name	Your choice
IPV4 CIDR Bock	Your choice. Refer to the CIDR block FAQ.

 ©2024 Palo Alto Networks, Inc.

Property	Value
Resource Group	Your Choice.

2. Select **Create VSwitch**.
   - Name the VSwitch **`Management`**.
   - Choose the **Zone**, specify an **IPv4 CIDR Block** that is a subset of the block you specified for the VPC, and specify a **Description**.
   - At the bottom, click **Add** to add another vSwitch (do not click **OK** until you have added all VSwitches).

   Refer to Create a VSwitch.
3. **Add** the Untrust VSwitch in the same manner.
4. **Add** the Trust VSwitch.
5. Click **OK**.

   View the VPC details and make any changes before you click **Complete**.

STEP 4 | Create security groups and security group rules.

- From the Alibaba Cloud Console home page, select **Elastic Compute Service** > **Networking & Security** > **Security Groups**.
- On the upper right, click **Create Security Group**.

1. Create the management security group.

   Refer to Create a security group to fill out the following fields.

Property	Value
Template	Customize
Security Group Name	Management
Security Group Type	Basic
Network Type	VPC
VPC	Select the VPC you created earlier.

 ©2024 Palo Alto Networks, Inc.

Property	Value
Resource Group	Your choice

- Complete the form and click **OK**.

  ECS console prompts you to create rules for this security group. This task describes some basic security group rules that allow you to bring up the VM-Series Firewall. You can create more rules to enforce your network security requirements.

2. Select **Create Rules Now** and create rules for HTTPS and SSH.

   Select the Inbound tab, and click **Add Security Group Rule**.

   - Create an Inbound rule to allow HTTPS in this security group. For example:

Property	Value
Rule Direction	Inbound
Action	Allow
Protocol Type	HTTPS (443)
Priority	100
Authorization Type	Refer to Add security group rules.
Authorization Object	

   - Click Add **Security Group Rule** to create an inbound rule to allow SSH on the management interface.

Property	Value
Rule Direction	Inbound
Action	Allow
Protocol Type	Customized TCP
Port Range	1/65535
Authorization Type	Refer to Add security group rules.

 ©2024 Palo Alto Networks, Inc.

Property	Value
Authorization Object	

Click **OK** and select **Back** to return to the Security Groups page.

3. Select **Create Security Group** and create the Untrust security group.

   When prompted, create a rule for the Untrust security group.

Property	Value
Rule Direction	Inbound
Action	Allow
Protocol Type	Custom TCP
Port Range	1/65535
Priority	100
Authorization Type	Refer to Add security group rules.
Authorization Object	

   Click **OK** and select **Back** to return to the Security Groups page.

4. Create the Trust security group.

   When prompted, click **Add Security Group Rule** and duplicate the Untrust rule.

   Continue to Create and Configure the VM-Series Firewall.

Create and Configure the VM-Series Firewall

This task uses the ECS console to create a VM-Series firewall instance with a minimum of three interfaces: management, untrust, and trust. An ECS instance supports a single NIC by default, and automatically attaches an Elastic Network Interface (ENI) to it. To support the VM-Series firewall, you must separately create the Untrust and Trust Elastic Network Interfaces (ENIs) and attach them to your instance.

STEP 1 | From the Alibaba Cloud console home page, select **Elastic Compute Service** > **Instances & Images** > **Instances**, and click **Create Instance** on the upper right.

STEP 2 | Select **Custom Launch**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Basic Configurations.

1. Fill in the following values. For example:

Property	Value
Billing Method	Subscription.
Region	Your choice. You can also select a Zone. The region you select must provide one of the required instance types.
Instance Type	One of the types in Alibaba Cloud Instance Type Recommendations for the VM-Series Firewall. You can use Type-based Selection to search for the instance type.
Image	Select **Marketplace Image** and search the Alibaba Marketplace for “VM-Series”. The image combines the OS and the VM-Series firewall.
Storage	Choose a disk type and specify 60 GB.
Snapshot	Your choice.

 ©2024 Palo Alto Networks, Inc.

Property	Value
Duration	Your choice.

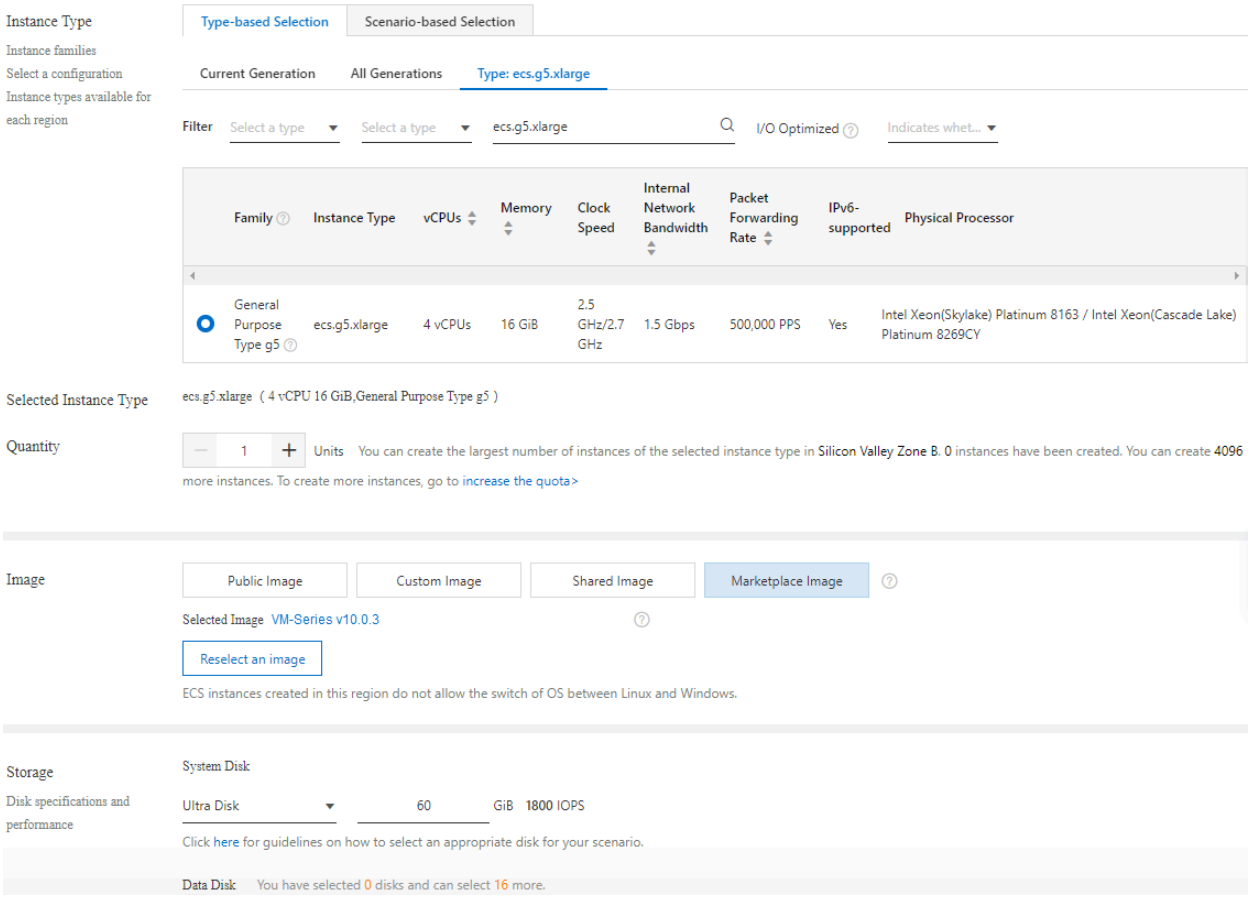

2. Select **Next: Networking**.

STEP 4 | On the Networking page, supply the following values.

1. Network (select VPC).
   - Choose the VPC you created in Create a VPC and Configure Networks.
   - Choose the Management VSwitch.
2. Public IP Address.

   If you do not have a public IP address, enable **Assign Public IP address** and the system will allocate one. If you must use a specific IP address, or an address in a specific range, you can request a custom IP address. Refer to the Elastic IP Address User Guide.
3. Security Group.

   Select the Management security group.
4. Elastic Network Interface.

   The Management interface is already attached to eth0.
5. Select **Next: System Configurations**.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | On the System Configurations page, fill in the following values.

1. Logon Credentials: Select **Key Pair**.

   Password authentication is not supported.

2. Name the VM-Series firewall instance and supply a Host name.

   Make any corrections.

   Select **Preview** to view your settings thus far.

3. Following **Advanced (based or instance RAM roles or cloud-init)** click **Show**.

   - The RAM role is optional.
   - In the User Data field, enter basic bootstrap information as key-value pairs separated by newlines. See Enter a Basic Configuration as User Data (Public Clouds). For example, enter the following in the **User Data** field.

```
type=dhcp-client
hostname=Ca-FW-DC1
vm-auth-key=7550362253****
panorama-server=10.*.*.20
panorama-server-2=10.*.*.21
tplname=FINANCE_TG4
dgname=finance_dg
op-cmd-dpdk-pkt-io=on
dhcp-send-hostname=yes
dhcp-send-client-id=yes
dhcp-accept-server-hostname=yes
dhcp-accept-server-domain=yes
authcodes=I7115398
vm-series-auto-registration-pin-id=abcdefgh1234****
vm-series-auto-registration-pin-value=zyxwvut-0987****
```

   op-command-modes *(mgmt-interface-swap and jumbo frame) are not supported for Alibaba Cloud.*

   op-cmd-dpdk-pkt-io=on *supports DPDK. If you want to specify PacketMMAP, specify op-cmd-dpdk-pkt-io=off*

   Grouping is Optional. Select **Preview** to view the configuration before ordering.

STEP 6 | View the terms of service, and select **Create Order** to create the VM-Series firewall instance.

View the purchase order and select **Subscribe**.

STEP 7 | From the console home page, choose > **Elastic Compute Service** > **Networks and Security** > **ENIs** and select **Create ENI** in the top right corner. Create elastic network interfaces for the Untrust and Trust interfaces.

1. Create the Untrust ENI.

   In the **Actions** column, select **Bind to Instance** and select the instance you just created.

2. Create the Trust ENI and bind it to the instance.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Allocate Elastic IP (EIP) addresses.

Allocate EIP addresses for the VM-Series firewall Management interface and the Untrust network interface. In this example the Trust interface is not exposed to the internet, so you don't need a third IP address.

If you already have two EIPs, go to the next step.

1. Associate an EIP with the VM-Series firewall Management interface.
2. Associate an EIP with the VM-Series firewall Untrust network interface.

   The second interface you attach is assigned to network interface 1 on the VM-Series firewall.

STEP 9 | Restart your instance to attach the new network interfaces.

On the Instances list, select your instance, select **Manage**, and select **Restart** on the upper right.

STEP 10 | SSH in to the VM-Series firewall with the security key and set the admin password:

```
developer1$ ssh -i dev1-vpc1.pem admin@18.***.145.153
Welcome admin.

admin> configure
Entering configuration mode
[edit]
admin# set mgt-config users admin password
Enter password:<password>
Confirm password:<password>
[edit]
admin# commit
```

STEP 11 | Access the VM-Series firewall web interface.

Open a web browser and enter the EIP for the management interface.

Secure North-South Traffic on Alibaba Cloud

After creating a VPC, you can create VSwitches to segment your virtual private network into subnets. This sample features a VPC with CIDR 192.168.0.0/16; you can enter your own values. Four VSwitches create four subnets.

VSwitch Name	Interface	Sample CIDR
mgmt	eth0	192.168.0.0/24
untrust	eth1	192.168.1.0/24
web	eth2	192.168.2.0/24
db	eth3	192.168.3.0/24

 ©2024 Palo Alto Networks, Inc.

In the following diagram, the VM-Series firewall connects to two trusted subnets, web and db. Inbound traffic is initiated when an external client accesses the VM-Series firewall's Untrust interface. The firewall inspects the traffic and sends it to an application. For example, the firewall sends traffic to a Web server through the Trust interface. The traffic returning from the Web server must hit the VM-Series firewall's Trust interface. The firewall inspects the return traffic flow, and sends it out through the Untrust interface.

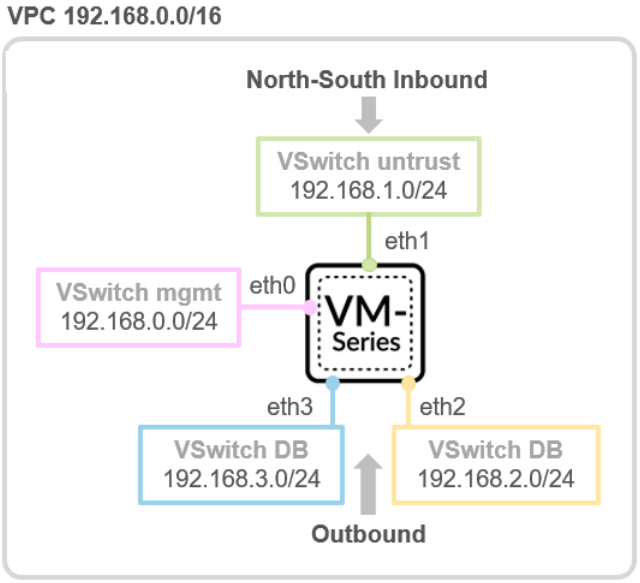

To secure inbound traffic, both DNAT and SNAT must be configured on the firewall.

STEP 1 | Create NAT rules for inbound traffic.

Here's a sample of the NAT rules for inbound traffic protection.

```
            <nat>
               <rules>
                 <entry name="inbound_web">
                   <source-translation>
                     <dynamic-ip-and-port>
                       <interface-address>
                         <interface>ethernet1/2</interface>
                       </interface-address>
                     </dynamic-ip-and-port>
                   </source-translation>
                   <destination-translation>
                     <translated-address>web_server</translated-
address>
                   </destination-translation>
                   <to>
                     <member>untrust</member>
                   </to>
                   <from>
                     <member>any</member>
                   </from>
                   <source>
                    <member>any</member>
                   </source>
```

 ©2024 Palo Alto Networks, Inc.

```
                  <destination>
                    <member>fw_untrust</member>
                  </destination>
                  <service>any</service>
                  <to-interface>ethernet1/1</to-interface>
                </entry>
              </rules>
            </nat>

          <address>
              <entry name="fw_untrust">
                <ip-netmask>192.168.1.4</ip-netmask>
              </entry>
              <entry name="fw_trust">
                <ip-netmask>192.168.2.201</ip-netmask>
              </entry>
              <entry name="web_server">
                <ip-netmask>192.168.2.203</ip-netmask>
              </entry>
          </address>
```

STEP 2 | Secure outbound traffic.

As shown in the diagram above, an application initiates the outbound traffic. For example, a web server must run **`yum install`** to update rpm packages. Typically the internet facing

 ©2024 Palo Alto Networks, Inc.

traffic within a VPC is routed to a NAT gateway (with an EIP attached). To secure outbound traffic, you must force outbound traffic to go through the VM-Series firewall.

1. Add a default gateway route in the VPC routing table with firewall IP in the subnet of the web server as the next hop.

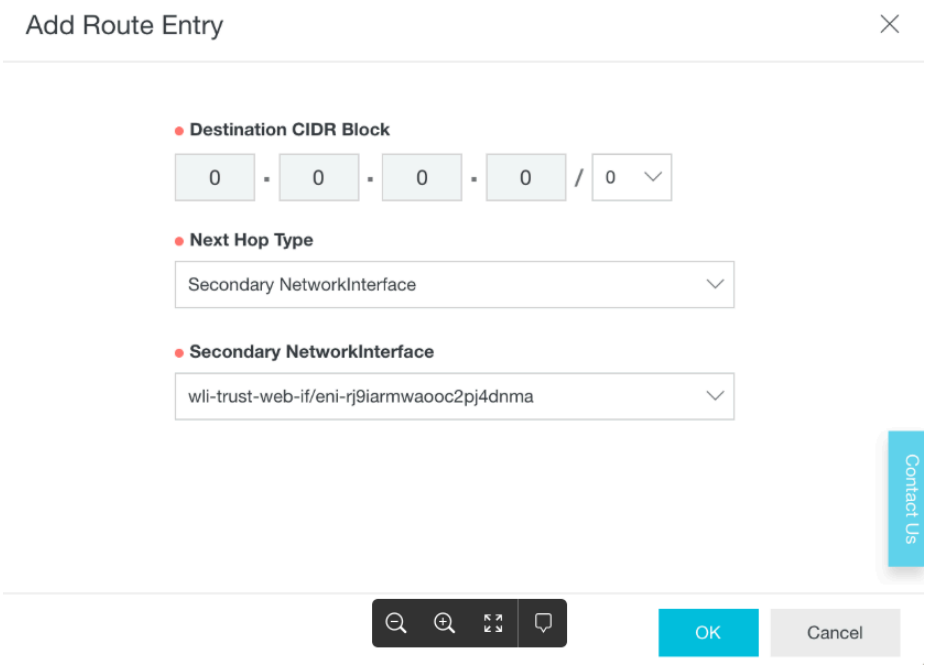

2. View your entry in the route table.

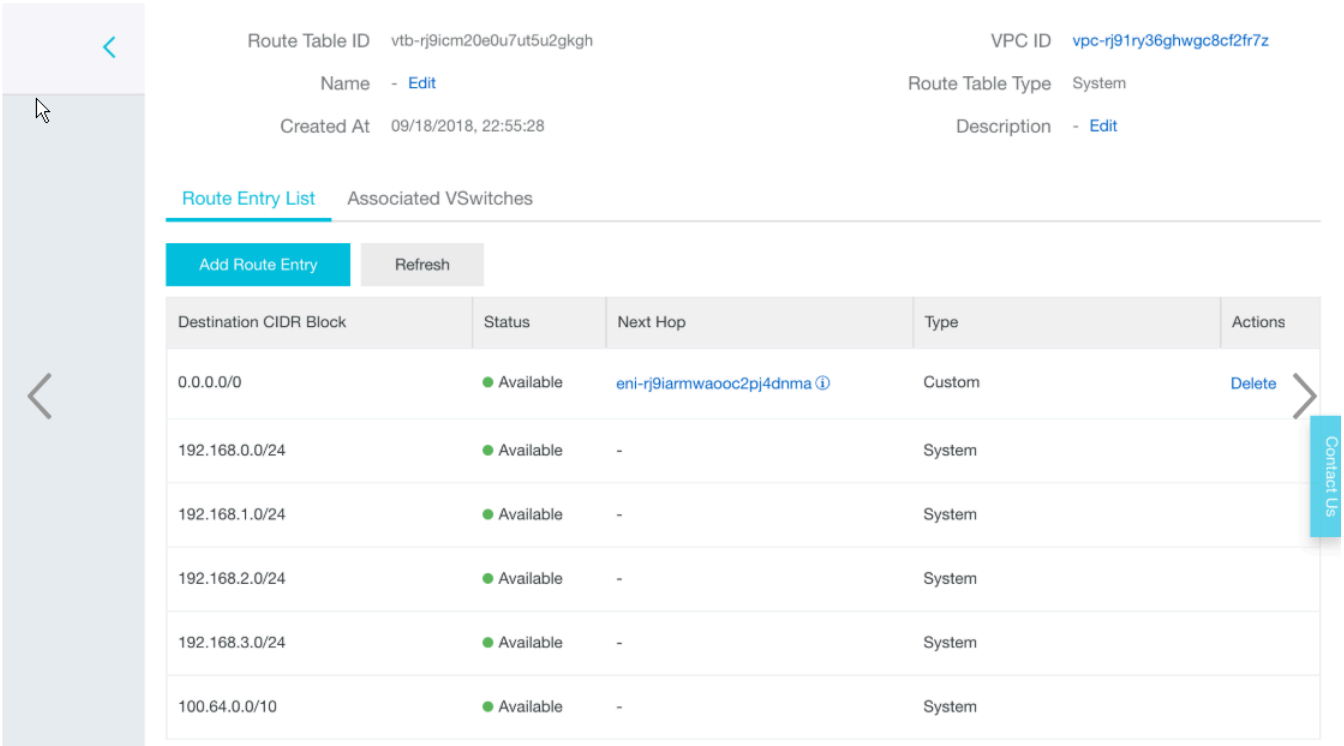

3. Configure SNAT rules using the Untrust interface IP to ensure traffic returning from the internet goes through the VM-Series firewall.

 ©2024 Palo Alto Networks, Inc.

Here's a sample SNAT configuration.

```
<nat>
  <rules>
    <entry name="outbound_web">
      <source-translation>
        <dynamic-ip-and-port>
          <interface-address>
            <interface>ethernet1/1</interface>
          </interface-address>
        </dynamic-ip-and-port>
      </source-translation>
      <to>
        <member>untrust</member>
      </to>
      <from>
        <member>trust</member>
      </from>
      <source>
        <member>any</member>
      </source>
      <destination>
        <member>any</member>
      </destination>
      <service>any</service>
      <to-interface>any</to-interface>
    </entry>
  </rules>
</nat>
```

Configure Load Balancing on Alibaba Cloud

On Alibaba Cloud, you can deploy the VM-Series firewall in a load balancer sandwich configuration where the firewall is deployed between a public network and a private network, as shown below.

 ©2024 Palo Alto Networks, Inc.

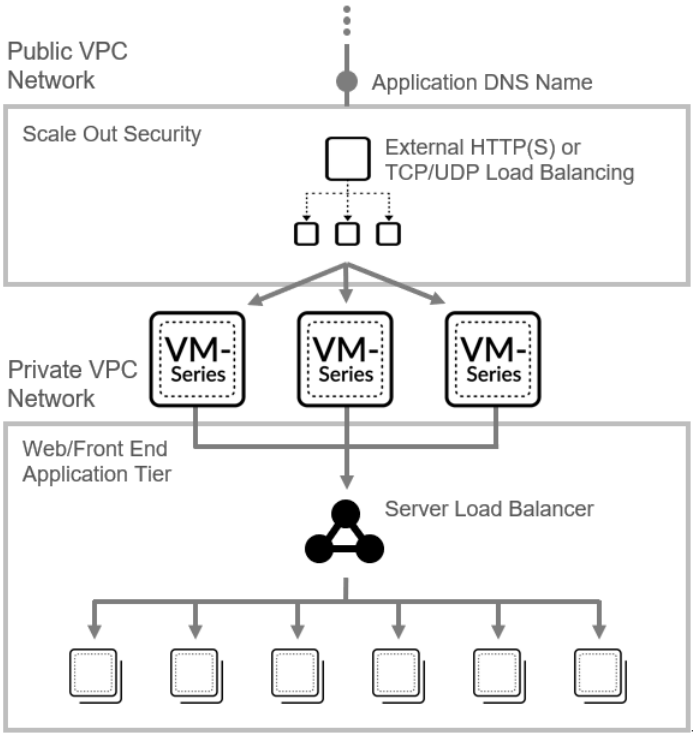

In Create a VPC and Configure Networks, you created Untrust and Trust ENIs and attached them to the VM-Series firewall instance as secondary ENIs.

When you use the console to add multiple backend servers to Alibaba Server Load Balancer (SLB), the SLB sends traffic to the primary ENI of the next-hop backend servers. Because the primary ENI is the management interface, traffic must go to the Untrust interface (a secondary ENI) for inspection.

To ensure that internet traffic goes to dataplane interfaces rather than the management interface, use the Alibaba CLI to attach the VM-Series firewall untrust ENIs to your SLB instance.

You must install the Aliyun command line interface to use the following CLI commands.

STEP 1 | Create the public and private VPCs for a load balancer sandwich configuration, and deploy the VM-Series firewalls.

The remaining steps are sample CLI commands you can adapt to your environment.

STEP 2 | Create the load balancer.

```
slb CreateLoadBalancer --RegionId us-west-1 --LoadBalancerName wli-
slb
--VpcId vpc-rj91ry36ghwgc8cf2fr7z --LoadBalancerSpec slb.s1.small
--AddressType internet --MasterZoneId us-west-1a
--SlaveZoneId us-west-1b
{
        "NetworkType": "classic",
        "LoadBalancerName": "**********",
        "Address": "**********,
        "ResourceGroupId": "rg-**********ofi",
```

 ©2024 Palo Alto Networks, Inc.

```
        "RequestId": "0B8BA2AA-E837-****-****-B82A8A1D5FBB",
        "AddressIPVersion": "ipv4",
        "LoadBalancerId": "lb-******************mvz",
        "VSwitchId": "",
        "VpcId": "vpc-*****************r7z"
}
```

STEP 3 | Add backend servers.

Use the CLI to add interfaces one at a time. The order in which you add the interfaces determines which NIC receives the interface.

```
aliyun slb AddBackendServers --LoadBalancerId lb-
******************mvz
--BackendServers
'[
        {
                "ServerId":"eni-******************bzw",
                "Type":"eni","Weight":"100"
        }
]'
```

STEP 4 | Create an HTTP Listener that performs a health check.

```
aliyun slb CreateLoadBalancerHTTPListener
--LoadBalancerId lb-******************mvz
--ListenerPort 80 --StickySession on
--HealthCheck on --HealthCheckURI '/'
```

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up a Firewall in Cisco ACI

Palo Alto Networks integrates as a service with Cisco Application-Centric Infrastructure (ACI). ACI is a software-defined networking (SDN) solution for easily deploying new workloads and network services. Using an SDN controller called the Cisco Application Policy Infrastructure Controller (APIC), you deploy the firewall service between Endpoint Groups (EPGs). EPGs act as a container for applications or application tiers. When you place a firewall between EPGs, security policy configured on the firewall secures the traffic between the EPGs. The APIC provides a single pane of glass for managing the network topology, network policies, and connectivity for the entire data center and supports inserting L4 - L7 devices, such as a hardware-based or VM-Series firewall. Panorama is required for centralized security management.

- Palo Alto Networks Firewall Integration with Cisco ACI
- Prepare Your ACI Environment for Integration
- Integrate the Firewall with Cisco ACI in Network Policy Mode
- Endpoint Monitoring in Cisco ACI

Palo Alto Networks Firewall Integration with Cisco ACI

Palo Alto Networks integration with Cisco ACI allows you to insert a firewall between EPGs as a Layer 4 to Layer 7 service. The firewall then secures the east-west traffic between the application tiers within those EPGs or north-south traffic between users and the applications.

The figure below shows an example of a physical ACI deployment that includes integrated Palo Alto Network firewalls. All the entities in the ACI Fabric are connected to leaf switches and those leaf switches are connected to larger spine switches. As users access the application, the ACI fabric moves the traffic to the correct destination. To secure the traffic between the application tiers, the network administrator inserts the Palo Alto Networks firewalls as L4 to L7 services between each EPG and creates a service graph to define what services the L4 to L7 device provides.

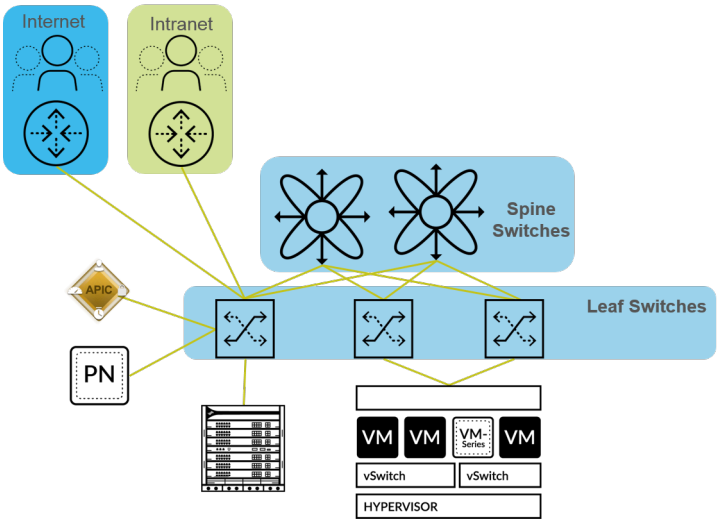

After the firewall services have been deployed, traffic now flows logically as shown below. Traffic to and from the end users and each tier in the application regardless of where or how each entity is physically connected to the network.

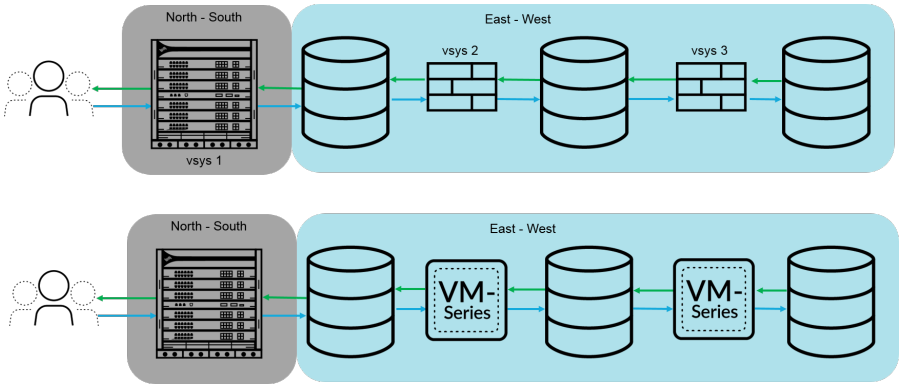

When the firewall is integrated with Cisco ACI, traffic is sent to the firewall with a policy-based redirect (PBR). Additionally, configuration of the firewall and configuration of the APIC are completely separate. Network policy mode does not rely on any other configuration integration between the firewall and the APIC, so it provides greater flexibility of configuration and deployment of the firewall.

 ©2024 Palo Alto Networks, Inc.

For east-west traffic, define a bridge domain and subnet in the ACI fabric for the firewall. Configure contracts between EPGs that send traffic to the firewall using a PBR. The PBR forwards traffic to the firewall based on policy containg the firewall's IP and MAC address. The firewall interfaces are always in Layer 3 mode and traffic is received and routed back to the ACI fabric. You can configure separate interfaces for consumer and provider connections or a single interface for ingress and egress traffic. The procedure in this document uses a single interface because it simplifies the integration; you do not need to configure as many interfaces, IP addresses, or VLANs. However, when using a single interface, you cannot uses zone information in defining security policy and you must modify the default intra-zone policy on the firewall to deny traffic.

For north-south traffic, you must use a dedicated policy called an L3Out. An L3Out contains the information required for the tenant to connect to external routing devices and access external networks. L3Out connections contain an external network EPG that represent the networks accessible through the L3Out policy. Just as the L3Out can group all external networks into a single EPG, you can use a vzAny object ACI to represent all EPGs in a VRF. Using a vzAny object simples the application of the outbound traffic contract because, whenever a new EPG is added to the VRF, the contract is automatically applied. In this scenario, the external network provides the contract and the vzAny object (all internal EPGs) consume it.

The following section provide additional details about components and concepts that make up the integration between the Next-Generation Firewall and Cisco ACI.

- Service Graph Templates
- Multi-Context Deployments

Service Graph Templates

Firewalls are deployed in Cisco ACI through service graphs. A service graph allows you to integrate Layer 4 - Layer 7 devices, such as a firewall, into the flow of traffic without the need for the L4-L7 device to be the default gateway for the servers in the ACI fabric.

Firewalls are represented in the ACI fabric as an L4-L7 device that you configure in the APIC as a device cluster. A single firewall or two firewalls deployed as an HA pair are configured as a device cluster. Each device cluster has one or more logical interfaces that describe the interface information of the device cluster and map the path of the member firewall with a VLAN from the physical or virtual machine monitor (VMM) domain.

Service graph templates define the firewall device cluster that you insert into the traffic flow between EPGs. Additionally, the service graph template defines the how the firewall is integrated and the logical interfaces that are assigned to the consumer and provider EPGs. After creating your service graph template, you assign it to EPGs and contracts. Because the service graph template is not tied to a specific EPG or contract, you can reuse it between multiple EPGs. The APIC then deploys the service graph template by connecting it to the bridge domain between EPGs.

Multi-Context Deployments

Cisco ACI integration supports physical firewalls divided into contexts that are managed by ACI as individual firewalls. On the firewall, these contexts are the virtual systems (vsys) on the firewalls and each firewall is licensed to support a certain number of vsys instances. When deploying a multi-vsys firewall in ACI, you must configure a chassis manager in the tenant and assign it to the firewall service.

 ©2024 Palo Alto Networks, Inc.

Prepare Your ACI Environment for Integration

Before you can integrate the firewall with a device package, you must complete the following steps to prepare your Cisco ACI environment.

STEP 1 | Deploy Panorama.

STEP 2 | Deploy the firewall.

- Physical Firewall—Connect the firewall's out-of-band management port to one leaf switch port and connect at least one firewall data interface to the switch. Firewall interfaces on a physical firewall are configured with VLANs to ensure connectivity to the correct networks. Deploy the firewall according to the platform-specific installation guide.
- VM-Series Firewall—When configuring the virtual hardware for the VM-Series firewall, set the port-group for the management interface. Each VM-Series firewall connected to the network requires its own virtual NIC. Deploy the VM-Series firewall based on your hypervisor.

STEP 3 | Configure the management IP address on each firewall and Panorama.

Perform initial configuration on:

- Hardware-based firewall
- VM-Series firewall
- Panorama

STEP 4 | Add your firewall(s) to Panorama as a managed device.

STEP 5 | Install feature licenses on your firewall(s).

- Register and activate licenses on your physical firewall.
- Register and Activate VM-Series Model Licenseson your VM-Series firewall.
- Manage firewall licenses using Panorama.

STEP 6 | Establish Cisco ACI fabric and management connectivity.

As part of this configuration, create a physical domain and VLAN namespace. Ensure that data interfaces of any physical firewalls are part of the physical domain.

STEP 7 | Create a Cisco ACI VMM domain profile.

If you are using virtual machines or the VM-Series firewall, create a virtual machine monitor (VMM) domain profile for the VMware vSphere environment. The VMM domain specifies the connectivity policy between vSphere and the ACI fabric.

 ©2024 Palo Alto Networks, Inc.

Integrate the Firewall with Cisco ACI in Network Policy Mode

In network policy mode, you integrate a pair of firewalls in high availability (HA) into the east-west or north-south traffic by using a policy-based redirect to a single logical HA interface. The firewall and ACI fabric are configured separately and address objects on the firewall are mapped to EPGs in the ACI fabric.

You can use network policy mode to deploy a Palo Alto Networks firewall to secure east-west or north-south traffic.

- Deploy the Firewall to Secure East-West Traffic in Network Policy Mode
- Deploy the Firewall to Secure North-South Traffic in Network Policy Mode

Deploy the Firewall to Secure East-West Traffic in Network Policy Mode

The following procedure describes how to deploy a Palo Alto Networks firewall to secure east-west traffic in the your Cisco ACI environment using unmanaged mode with policy-based redirect. This procedure assumes that you have completed the following:

- Firewalls are operational and connected to a leaf switch in your Cisco ACI environment. Additionally, the management interface of each firewall must be reachable by the APIC.
- Firewalls are deployed in active/passive HA mode. This procedure does not cover HA network setup and assumes you have completed this in advance.

To secure east-west traffic, define a bridge domain and subnet in the ACI fabric for the firewall. Configure contracts between EPGs that send traffic to the firewall using a PBR. The PBR forwards traffic to the firewall based on policy containing the firewall's IP and MAC address. The firewall interfaces are always in Layer 3 mode and traffic is received and routed back to the ACI fabric. You can configure separate interfaces for consumer and provider connections or a single interface for ingress and egress traffic. The procedure in this document uses a single interface because it simplifies the integration; you do not need to configure as many interfaces, IP addresses, or VLANs. However, when using a single interface, you cannot uses zone information in defining security policy and you must modify the default intra-zone policy on the firewall to deny traffic.

This procedure deploys the firewall in one-arm mode. In one-arm mode, the traffic enters and exits the firewall through a single interface. This common firewall interface is used for both consumer and provider interfaces in the service graph template. Using a single interface simplifies integration with the firewall by reducing the number IP addresses, VLANs, and interfaces that you must configure. However, a one-arm deployment model is intrazone, so you cannot use zone information to define security policy.

On the firewall:

- Create a Virtual Router and Security Zone
- Configure the Network Interfaces
- Configure a Static Default Route
- Create Address Objects for the EPGs

 ©2024 Palo Alto Networks, Inc.

- Create Security Policy Rules

On the Cisco APIC:

- Create a VLAN Pool and Domain
- Configure an Interface Policy for LLDP and LACP for East-West Traffic
- Establish the Connection Between the Firewall and ACI Fabric
- Create a VRF and Bridge Domain
- Create an L4-L7 Device
- Create a Policy-Based Redirect
- Create and Apply a Service Graph Template

Create a Virtual Router and Security Zone

Configure a virtual router and zone on the firewall for each VRF in the tenant.

STEP 1 | Log in to the firewall.

STEP 2 | Select **Network** > **Virtual Routers** and click **Add**.

STEP 3 | Give the virtual router a descriptive **Name**.

STEP 4 | Click **OK**.

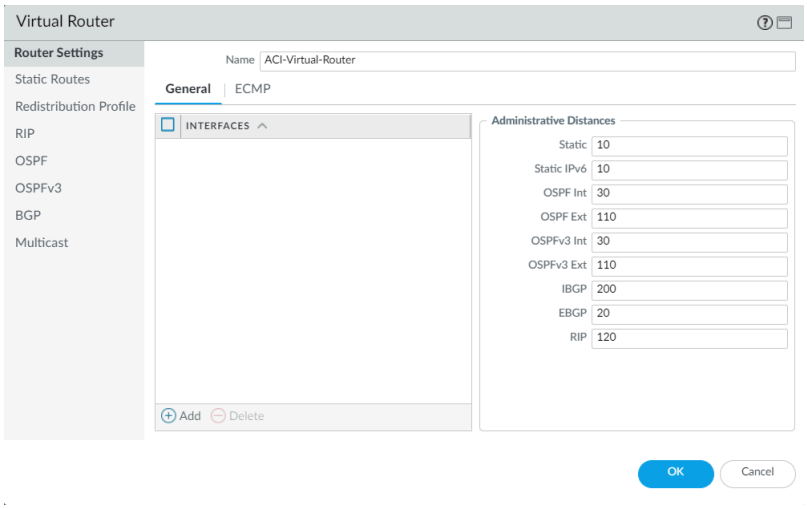

STEP 5 | Select **Network** > **Zones** and click **Add**.

STEP 6 | Give the zone a descriptive **Name**.

STEP 7 | Choose Layer 3 from the **Type** drop-down.

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Click **OK**.

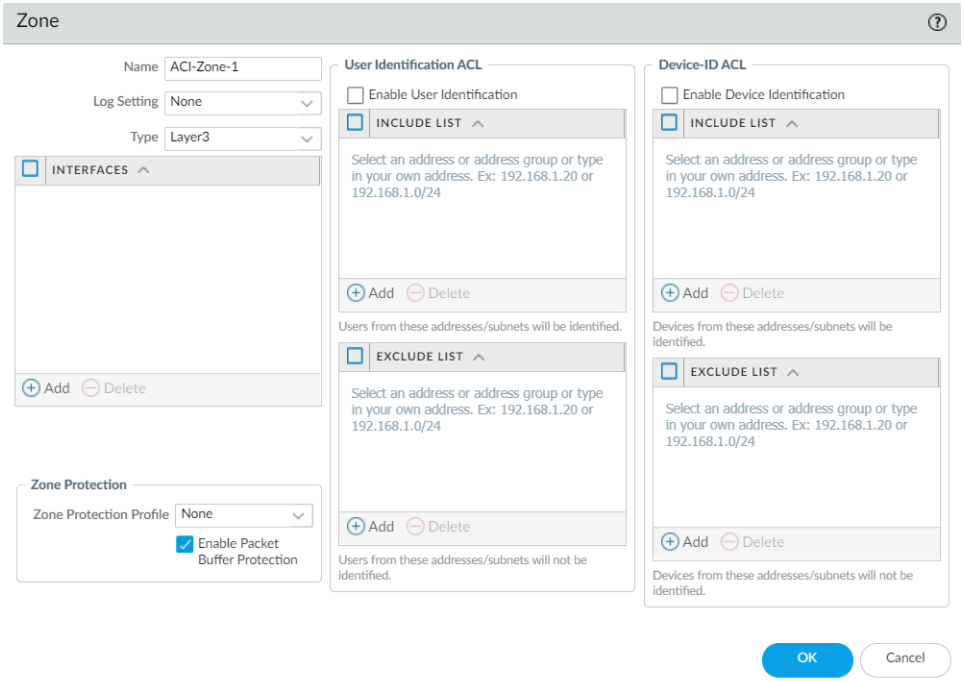

STEP 9 | **Commit** your changes.

Configure the Network Interfaces

Configure the Ethernet interfaces that connect the firewall to the ACI leaf switches. The VLAN ID number used in this configuration should be a member of the VLAN pool assigned to the firewalls in ACI.

The VM-Series firewall does not support aggregate Ethernet groups.

STEP 1 | Select **Network** > **Interfaces** > **Ethernet** and click **Add Aggregate Group**.

STEP 2 | Enter a number for the aggregate group in the second **Interface Name** field.

STEP 3 | Select Layer 3 from the **Interface Type** drop-down.

STEP 4 | Select the **LACP** tab and click **Enable LACP**.

STEP 5 | Select **Fast** as the **Transmission Rate**.

STEP 6 | Under High Availability Options, select **Enable in HA Passive State**.

*Do not select **Same System MAC Address for Active-Passive HA**. This option makes the firewall pair appear as a single device to the switch, so traffic will flow to both firewalls instead of just the active firewall.*

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Click **OK**.

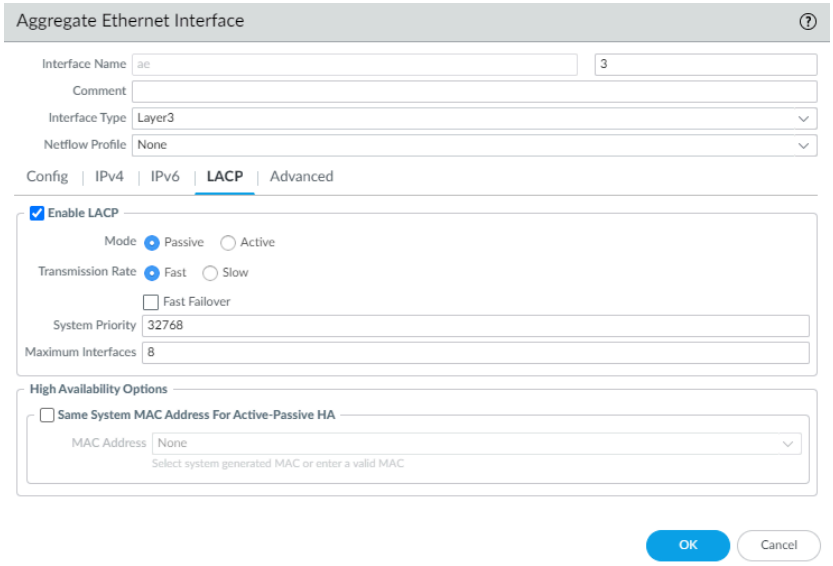

STEP 8 | Click on the name of an Ethernet interface to configure it and add it to the aggregate group.

1. Select **Aggregate Ethernet** from the Interface Type drop-down.
2. Select the interface you defined in the aggregate Ethernet group configuration.
3. Click **OK**.
4. Repeat this step for each other member interface of the aggregate Ethernet group.

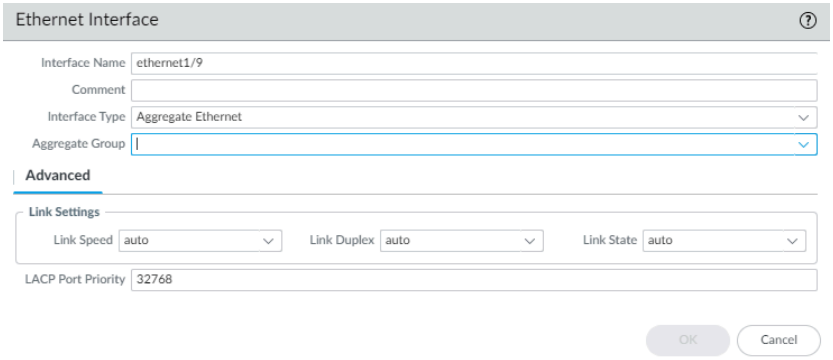

STEP 9 | Add a subinterface on the aggregate Ethernet interface for the tenant and VRF.

1. Select the row of your aggregate Ethernet group and click **Add Subinterface**.
2. In the second **Interface Name** field, enter a numerical suffix to identify the subinterface.
3. In the **Tag** field, enter the VLAN tag of the subinterface.
4. Select the virtual router you configured previously from the **Virtual Router** drop-down.
5. Select the zone you configured previously from the **Zone** drop-down.
6. Select the **IPv4** tab.
7. Select the **Static** Type.
8. Click **Add** and enter the subinterface IP address and network mask in CIDR notation.
9. Click **OK**.

 ©2024 Palo Alto Networks, Inc.

Configure a Static Default Route

Configure a static default route to direct traffic from the Ethernet subinterfaces to the subnet router.

STEP 1 | Select **Network** > **Virtual Routers** and click on the virtual router you created previously in this procedure.

STEP 2 | Select **Static Routes** > **IPv4** and click **Add**.

STEP 3 | Enter a descriptive **Name**.

STEP 4 | Enter 0.0.0.0/0 in the **Destination** field.

STEP 5 | From the **Interface** drop-down, select the aggregate Ethernet group you created previously in this procedure.

STEP 6 | Select IP Address from the **Next Hop** drop-down and enter the IP address of the next hop router.

STEP 7 | Click **OK**.

STEP 8 | Click **OK** again.

STEP 9 | **Commit** your changes.

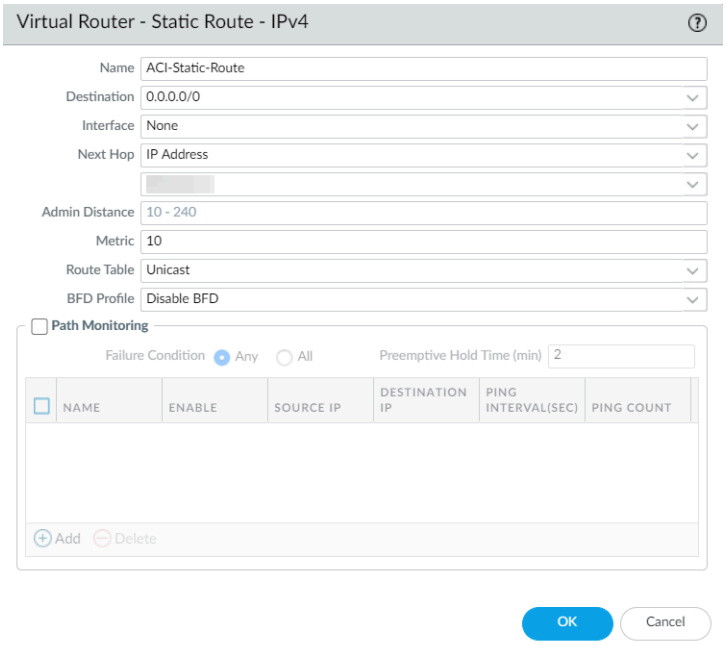

Create Address Objects for the EPGs

You must define address objects and map them to endpoint groups (EPGs) to be used in security policy. Address groups are the best way map security groups to a group of servers using an endpoint IP address range. Create one address object for each of your EPGs.

STEP 1 | Select **Objects** > **Address** and click **Add**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Enter a descriptive name for your address object.

STEP 3 | Select IP Netmask from the **Type** drop-down.

STEP 4 | Enter the IP Netmask.

STEP 5 | Click **OK**.

STEP 6 | Repeat this process for each EPG.

STEP 7 | **Commit** your changes.

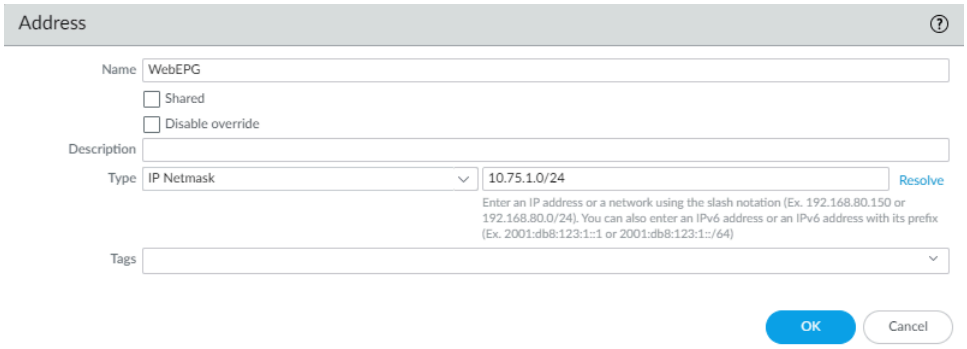

Create Security Policy Rules

Create security policy rules to control the traffic moving between your EPGs. By default, the firewall allows all intrazone traffic. Therefore, because the EPGs are in the same zone, all between those EPGs is allowed. Before creating a new rules, you will change the default intrazone rule from allow to deny.

STEP 1 | Select **Policies** > **Security**.

STEP 2 | Click on intrazone-default to highlight the row and click **Override**.

STEP 3 | Select the **Action** tab.

STEP 4 | Select Deny from the **Action** drop-down.

STEP 5 | Click **OK**.

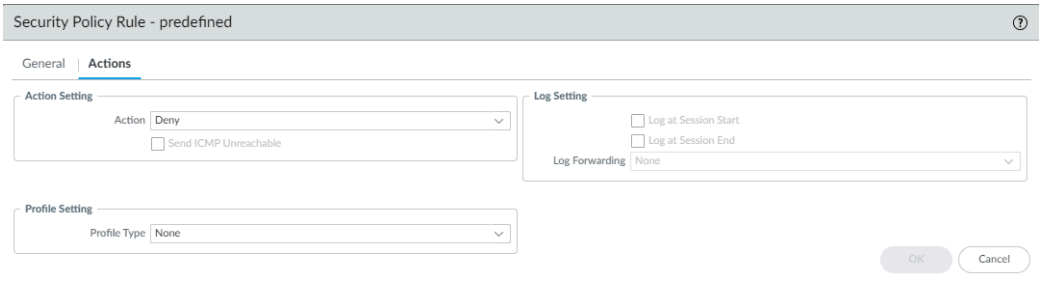

STEP 6 | Configure additional security policy rules based on your needs using the address objects and zone you created for your EPG.

 ©2024 Palo Alto Networks, Inc.

Create a VLAN Pool and Domain

Configure the VLAN pool that will be used to allocate VLANs to the firewall when you attach interfaces to the ACI infrastructure for EPGs. The firewall's VLAN pull should have a static VLAN range.

Configure a dedicated domain for the firewall. A domain for the firewall is required to map the VLANs to the EPGs. Create a physical domain for a physical firewall and create a VMM domain for a VM-Series firewall.

STEP 1 | Create a VLAN pool.

1. Log in to your APIC.
2. Select **Fabric** > **Access Policies** > **Pools** > **VLAN**.
3. Right-click **VLAN** and select **Create VLAN Pool**.
4. Enter a descriptive **Name** for your VLAN pool.
5. Select **Dynamic Allocation** for Allocation Mode.
6. Click the plus (+) button to the right of **Encap Blocks**.
7. Enter your VLAN range in the **VLAN Range** field.
8. Select **Static Allocation** form the Allocation Mode drop-down.
9. Click **OK**.
10. Click **Submit**.

STEP 2 | (Physical firewall only) Create a physical domain.

1. Select **Fabric** > **Access Policies** > **Physical and External Domains** > **Physical Domains**.
2. Right-click **Physical Domain** and select **Create Physical Domain**.
3. Enter a descriptive **Name** for your physical domain.
4. Select the VLAN pool you created in the previous procedure from the VLAN Pool list.
5. Click **Submit**.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | (VM-Series firewall only) Create a VMM domain.

1. Select **Virtual Networking** > **VMM Domains** > **VMware**.
2. Right-click **VMware** and select **Create vCenter Domain**.
3. Enter a descriptive **Name** for your VMM domain.
4. Select **VMware vSphere Distributed Switch** from the **Virtual Switch** drop-down.
5. Select **VLAN** from the **Encapsulation** drop-down.
6. Select your VLAN pool from the **VLAN Pool** drop-down.
7. Click the plus (+) button to the right of **vCenter Credentials**.
8. Enter a descriptive **Profile Name** and your vCenter login information.
9. Click the plus (+) button to the right of **vCenter**.
10. Enter a descriptive **Name**.
11. Select vCenter from the Type drop-down.
12. Enter your vCenter IP address under **IP/Hostname**.
13. Select the vCenter Credentials profile you just created from the **Associated Credential** drop-down.
14. Click **Submit**.

Configure an Interface Policy for LLDP and LACP for East-West Traffic

Create policy that enables LLDP and LACP on the ACI interfaces that connect to your firewall.

LLDP is necessary for forwarding to work correctly in the ACI environment; ACI does not deploy a subnet router interface on a leaf switch unless it detects an endpoint on the switch that requires one. LLDP helps determine if a subnet router interface is required.

LACP provides greater resiliency and recovery speed on a link failure.

STEP 1 | Create an LLDP Interface Policy.

1. Select **Fabric** > **Access Policies** > **Interface Policies** > **Policies** > **LLDP Interface**.
2. Right-click on **LLDP Interface** and select **Create LLDP Interface Policy**.
3. Enter a descriptive **Name** for your LLDP interface policy.
4. Select **Enabled** for **Receive State**.
5. Select **Enabled** for **Transmit State**.
6. Click **Submit**.

STEP 2 | Create a Port Channel policy to enable LACP.

1. Select **Fabric** > **Access Policies** > **Interface Policies** > **Policies** > **Port Channel**.
2. Right-click on **Port Channel** and select **Create Port Channel Policy**.
3. Enter a descriptive **Name** for your port channel policy.
4. Select **LACP Active** from the **Mode** drop-down.
5. Click **Submit**.

 ©2024 Palo Alto Networks, Inc.

Establish the Connection Between the Firewall and ACI Fabric

Attach your firewall to the leaf switch through a VPC connection using the Ethernet interface (or aggregate Ethernet group) you configured on your firewall earlier in this procedure. Connect the interface or interfaces to the same ports on the leaf switches.

STEP 1 | Select **Fabric** > **Access Policies** > **Quick Start**.

STEP 2 | Click **Configure an interface, PC, and VPC**.

STEP 3 | Click the green and white plus (+).

STEP 4 | Select the leaf switch or switches to which you firewall is connected from the **Switches** drop-down.

STEP 5 | Click the green and white plus (+).

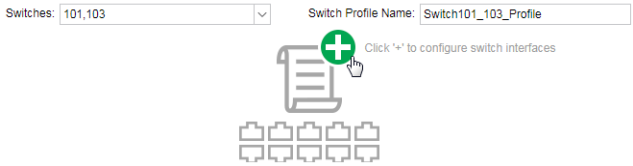

STEP 6 | Select VPC as the **Interface Type**.

STEP 7 | In the **Interfaces** field, enter the number of the interface your firewall uses to connect to the leaf switch.

STEP 8 | Enter a descriptive name into the **Interface Selector Name** field.

STEP 9 | Select **LLDP-Enabled** from the **LLDP Policy** drop-down.

STEP 10 | Select **LACP Active** from the **Port Channel Policy** drop-down.

STEP 11 | Select **Bare Metal** for a physical firewall or **ESX Hosts** for the VM-Series from the **Attached Device Type** drop-down.

STEP 12 | Select **Choose One** for **Domain**.

STEP 13 | Select the physical domain or VMM domain you created previously in this procedure from the **Domain** drop-down.

STEP 14 | Click **Save**.

STEP 15 | Click **Save** and then **Submit**.

STEP 16 | Repeat this procedure for the second firewall in your HA pair.

 ©2024 Palo Alto Networks, Inc.

Create a VRF and Bridge Domain

A tenant requires a VRF for all bridge domains and subnets. In this example, you will create a single, common VRF for the firewall and endpoints. Then configure a dedicated bridge domain for your firewall and disable dataplane learning. Disabling dataplane learning is required to use Policy Based Redirect in a bridge domain.

STEP 1 | Create a VRF.

1. On the **Tenants** tab, double-click on the name of your tenant.
2. Select **Networking** > **VRFs**.
3. Right-click **VRFs** and select **Create VRF**.
4. Enter a descriptive **Name** for your VRF.
5. Clear the **Create A Bridge Domain** check box.
6. Click **Finish**.

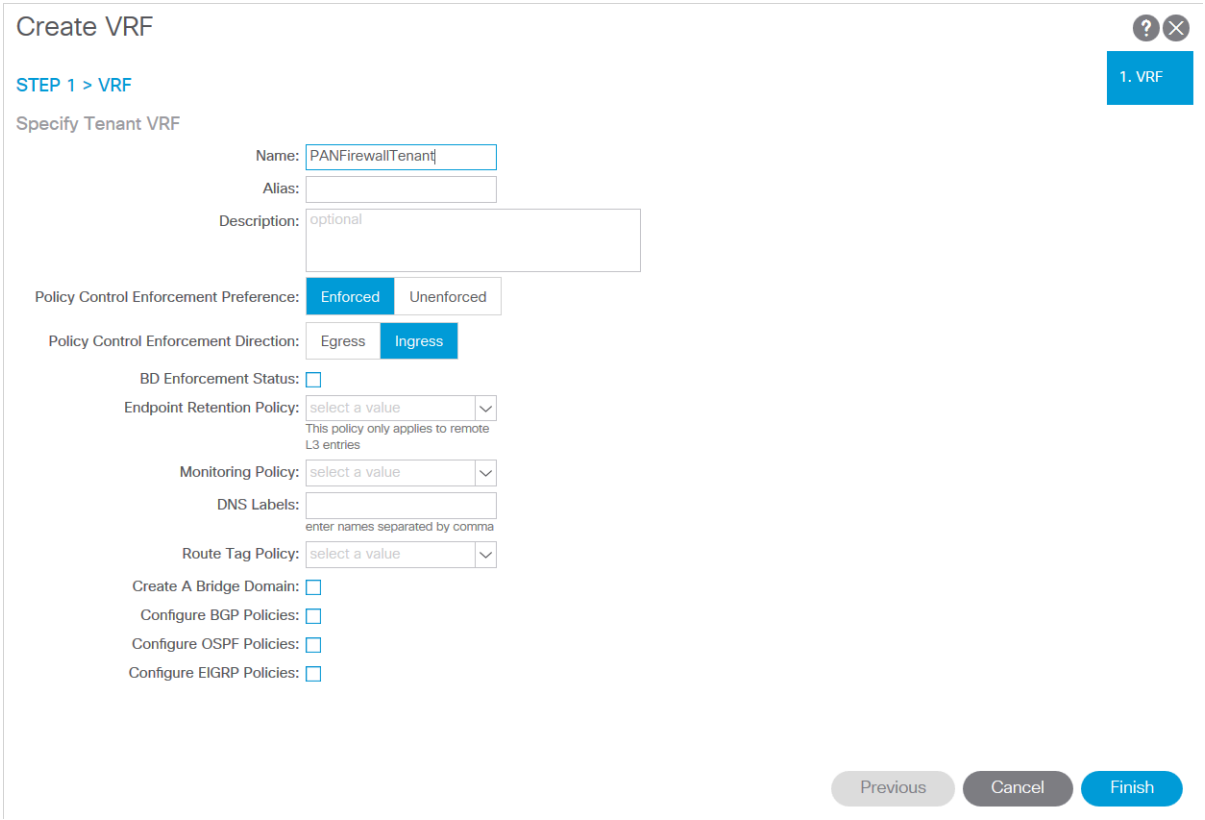

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Create a bridge domain for the firewall.

1. On the **Tenants** tab, double-click on the name of your tenant.
2. Select **Networking** > **Bridge Domains**.
3. Right-click **Bridge Domains** and select **Create Bridge Domain**.
4. Enter a descriptive **Name** for your bridge domain.
5. Select the VRF you created in the previous procedure from the **VRF** drop-down.
6. Click **Next**.

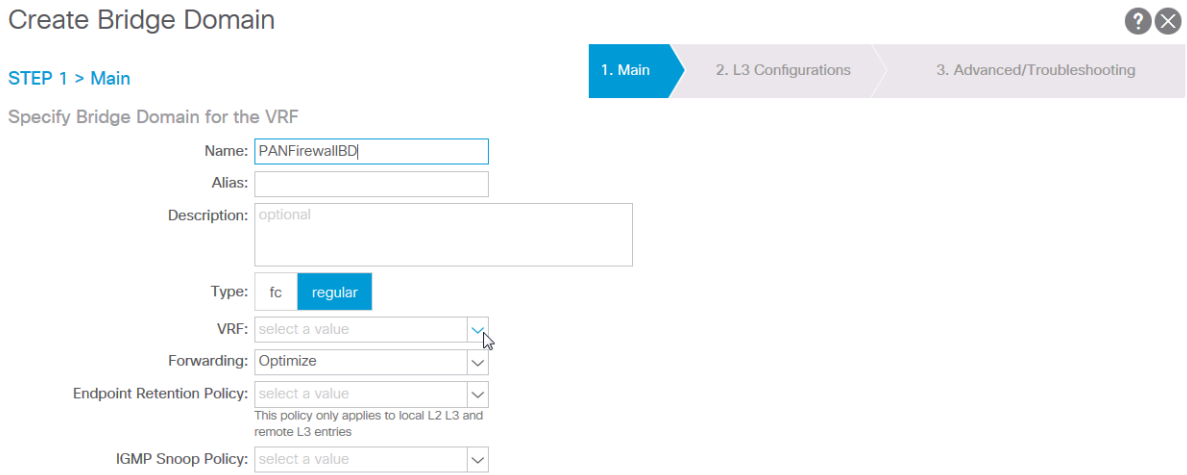

Create an L4-L7 Device

You must define the firewall as an L4-L7 device in the APIC so ACI can insert it into the traffic flow. You configure L4-L7 devices in the APIC as a device cluster, which is a construct that represents a single firewall or a firewall HA pair acting as a single device. Device clusters have one or more logical interfaces, which define the path of the member firewalls with a VLAN from the physical domain.

STEP 1 | On the **Tenants** tab, double-click on the name of your tenant.

STEP 2 | Select **Services** > **L4-L7** > **Devices**.

STEP 3 | Right-click **Devices** and select **Create L4-L7 Device**.

STEP 4 | Clear the **Managed** check box.

STEP 5 | Enter a descriptive **Name** for your L4-L7 Device.

STEP 6 | Select **Firewall** from the **Service Type** drop-down.

STEP 7 | Select **Physical** for a physical firewall or **Virtual** for a VM-Series firewall from the **Device Type** drop-down.

STEP 8 | Select the physical or VMM domain you created previously from the **Domain** drop-down.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Select HA Node for **View**.

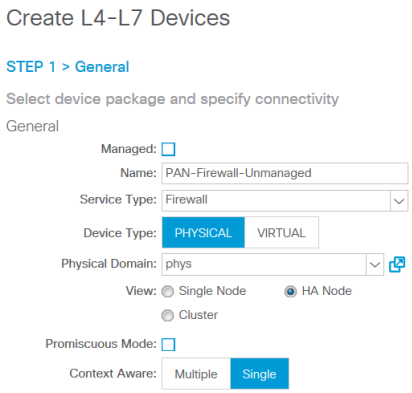

STEP 10 | Under **Device 1**, click the plus (+) icon to the right of **Device Interfaces**.

STEP 11 | Enter a descriptive **Name** for this interface.

STEP 12 | Under **Path**, select the path to the primary firewall in your HA pair.

STEP 13 | Click **Update**.

STEP 14 | Under **Device 2**, click the plus (+) icon to the right of **Device Interfaces**.

STEP 15 | Enter a descriptive **Name** for this interface.

STEP 16 | Under **Path**, select the path to the secondary firewall in your HA pair.

STEP 17 | Click **Update**.

STEP 18 | Under **Cluster**, click the plus (+) icon to the right of **Cluster Interfaces**.

STEP 19 | Enter a descriptive **Name** for the cluster.

STEP 20 | Select the two interfaces you configured above from the list under **Concrete Interfaces**. The APIC requires that you configure two interfaces. However, because there is only one connection between the firewall and the ACI fabric, only one of the interfaces is used.

STEP 21 | Under **Encap**, enter a VLAN from the from the static VLAN pool you created earlier. Traffic will be redirected to the firewall on the VLAN assigned here.

STEP 22 | Click **Update**.

 ©2024 Palo Alto Networks, Inc.

STEP 23 | Click **Finish**.

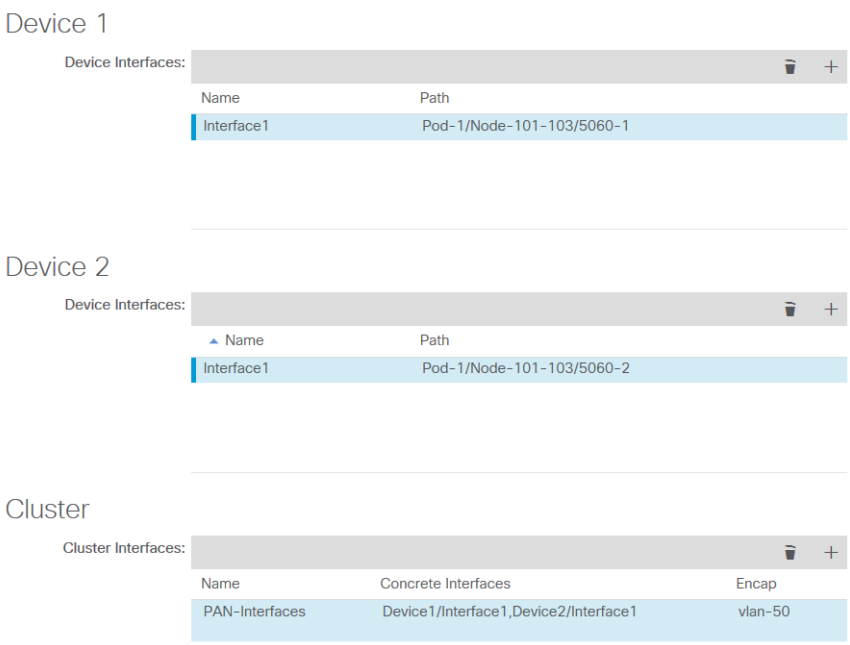

Create a Policy-Based Redirect

Configure the policy based redirect that sends the traffic between your EPGs to the firewall. Policy based redirect leverages the MAC address of the interface on the firewall. Before configuring the PBR setting on the APIC, you must get the MAC address from the firewall.

STEP 1 | Get the MAC address of the firewall.

1. Log into the firewall CLI.
2. Use the **`command show interface all`** to display the MAC addresses of your configured interfaces.
3. Copy the MAC address of the interface that will receive the redirected traffic.

STEP 2 | Create the L4-L7 Policy-Based Redirect.

1. Log into the APIC.
2. On the **Tenants** tab, double-click on the name of your tenant.
3. Select **Policies** > **Protocol** > **L4-L7 Policy Based Redirect**.
4. Right-click **L4-L7 Policy Based Redirect** and select **Create L4-L7 Policy Based Redirect**.
5. Enter a descriptive **Name** for your Policy Based Redirect.
6. Click the plus (+) icon to the right of **Destinations**.
7. In the **IP** field, enter the IP address of the interface that will receive the redirected traffic.
8. In the **MAC** field, enter the MAC address that you copied from the firewall CLI.
9. Click **OK**.
10. Click **Submit**.

 ©2024 Palo Alto Networks, Inc.

Create and Apply a Service Graph Template

Create a service graph template that uses the device cluster representing the firewall in a policy-based redirect integration. After creating the service graph, you must apply it to EPGs to protect traffic. A contract and contract filter rules define the traffic that can be forwarded to the firewall.

STEP 1 | Create a service graph template.

1. On the **Tenants** tab, double-click on the name of your tenant.
2. Select **Services** > **L4-L7** > **L4-L7 Service Graph Templates**.
3. Right-click **L4-L7 Service Graph Template** and select **Create L4-L7 Service Graph Template**.
4. Enter a descriptive **Graph Name** for your service graph template.
5. Select **Create a New One** for **Graph Type**.
6. Click and drag the L4-L7 device you created in the previous procedure between the consumer and provider EPGs.
7. Select **Routed** for **Firewall**.
8. Select **Routed Redirect**.
9. Click **Submit**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Apply the service graph template.

1. On the **Tenants** tab, double-click on the name of your tenant.
2. Select **Services** > **L4-L7**.
3. In the **EPGs Information** pane, select your consumer and provider EPGs from the **Consumer EPG** and **Provider EPG** drop-downs.
4. Select **Create a New Contract**.
5. Enter a descriptive **Contract Name**.
6. Clear No Filter **(Allow All Traffic)**. Using this option is not recommended. To allow all traffic between the EPGs to be redirected to the firewall, it is recommended that you create a filter to do this.
7. Click the plus (+) icon to the right of **Filter Entries**.
8. Create a rule (or rules) to define what traffic is allowed to pass between the EPGs and redirected to the firewall.
9. Click **Next**.

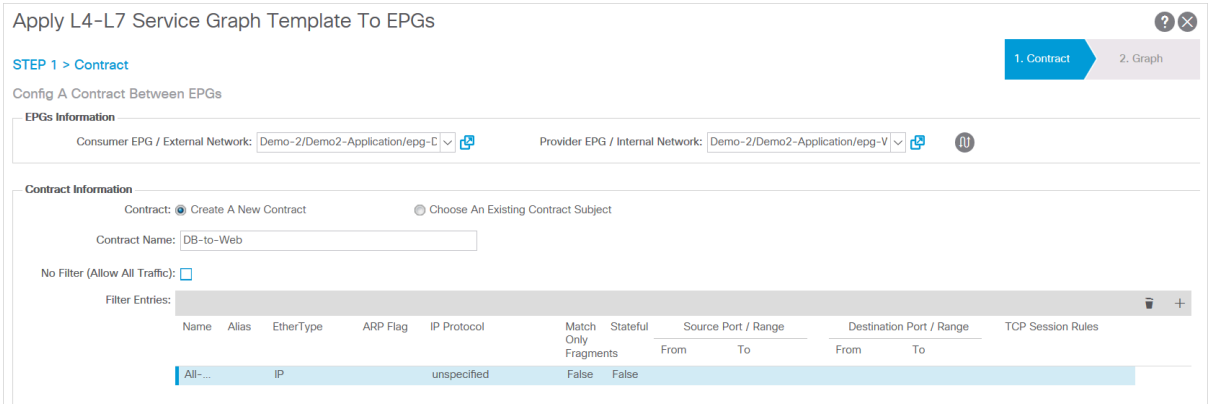

10. Select the service graph template you created in the previous procedure from the **Service Graph Template** drop-down.
11. In the consumer and provider pane, select the bridge domain containing your firewall from the **BD** drop-downs.
12. Select the policy based redirect you created previously from the **Redirect Policy** drop-downs.
13. Select the cluster interface you created with you L4-L7 device from the **Cluster Interface** drop-downs.

 ©2024 Palo Alto Networks, Inc.

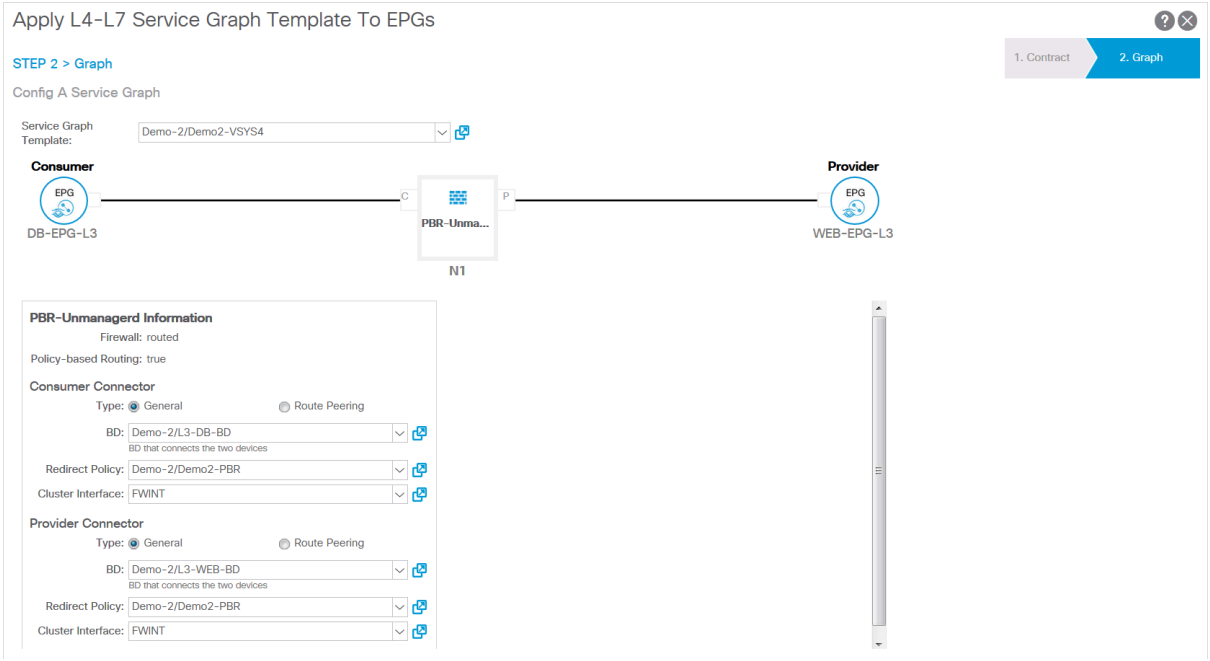

Deploy the Firewall to Secure North-South Traffic in Network Policy Mode

Use network policy mode to secure north-south traffic entering and exiting your data center using unmanaged mode with policy-based redirect. This procedure assumes that you have completed the following:

- Firewalls are operational and connected to a leaf switch in your Cisco ACI environment. Additionally, the management interface of each firewall must be reachable by the APIC.
- Firewalls are deployed in active/passive HA mode. This procedure does not cover HA network setup and assumes you have completed this in advance.

To establish external connectivity to networks outside of your ACI fabric, you must configure an L3Out. And L3Out is a dedicated policy that contains the parameters required to connect external routing devices to a tenant. Additionally, an L3Out contain an external EPG (called an external network in the APIC UI) that represents networks accessible through the L3Out. The external EPG is not dynamically populated and follows a zero-trust model, so you must define the networks in the EPG. To make configuration easier, you can configure a network of 0.0.0.0/0 to assign all networks to the external EPG.

To secure inbound traffic, connect your firewall or firewalls in an HA pair to your border-leaf switches. Border-leaf switches are leaf switches that provide Layer 3 connections to external routers. The firewalls peer with the border-leaf switches using the open shortest path first (OSPF) protocol that is configured on each leaf switch in the vPC pair and communicates with the firewalls using a switch virtual interface (SVI). On the firewall, you configure a virtual router dedicated to the interfaces that connect to your data center. Additionally, this procedure includes

For outbound traffic, the firewall advertises the external networks to the border-leaf switches using OSPF. Additionally, the external network EPG is configured to allow all networks advertised by the firewall into that EPG. You create a contract between a vzAny managed object and the external networks EPG to allow traffic from any EPG within the VRF to reach the external

 ©2024 Palo Alto Networks, Inc.

networks through the firewall. The vzAny managed object allows you to consolidate all EPGs in a VRF to one or more contracts instead of creating a separate contracts for each EPG. The EPGs collected in the vzAny managed object consume the contact provided by the external EPG.

Unlike in service manager mode, management if the ACI infrastructure and the firewalls is completed separately.

On the APIC—

- Create a VLAN Pool and External Routed Domain
- Configure an Interface Policy for LLDP and LACP for North-South Traffic
- Create an External Routed Network
- Configure Subnets to Advertise to the External Firewall
- Create an Outbound Contract
- Create an Inbound Web Contract
- Apply Outbound and Inbound Contracts to the EPGs

On the firewall—

- Create a Virtual Router and Security Zone for North-South Traffic
- Configure the Network Interfaces
- Configure Route Redistribution and OSPF
- Configure NAT for External Connections

Create a VLAN Pool and External Routed Domain

Create a VLAN pool to allocate VLANs to the firewall as you attach interfaces to the infrastructure to support the EPGs in your ACI fabric. You should use a static VLAN range for the firewall.

Additionally, you must create a physical domain to map the VLANs to the EPGs. The following procedure creates a physical domain dedicated to the firewall.

STEP 1 | Create a VLAN pool.

1. Log in to your APIC.
2. Select **Fabric** > **Access Policies** > **Pools** > **VLAN**.
3. Right-click **VLAN** and select **Create VLAN Pool**.
4. Enter a descriptive **Name** for your VLAN pool.
5. Select **Dynamic Allocation** for Allocation Mode.
6. Click the plus (+) button to the right of **Encap Blocks**.
7. Enter your VLAN range in the **VLAN Range** field.
8. Select **Static Allocation** form the Allocation Mode drop-down.
9. Click **OK**.
10. Click **Submit**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Create an external routed domain.

1. Select **Fabric** > **Access Policies** > **Physical and External Domains** > **External Domains**.
2. Right-click **External Routed Domain** and select **Create Layer 3 Domain**.
3. Enter a descriptive **Name** for your physical domain.
4. Select the VLAN pool you created in the previous procedure from the VLAN Pool list.
5. Click **Submit**.

Configure an Interface Policy for LLDP and LACP for North-South Traffic

Create policy that enables LLDP and LACP on the ACI interfaces that connect to your firewall.

LLDP is necessary for forwarding to work correctly in the ACI environment; ACI does not deploy a subnet router interface on a leaf switch unless it detects an endpoint on the switch that requires one. LLDP helps determine if a subnet router interface is required.

LACP provides greater resiliency and recovery speed on a link failure.

STEP 1 | Create an LLDP Interface Policy.

1. Select **Fabric** > **Access Policies** > **Interface Policies** > **Policies** > **LLDP Interface**.
2. Right-click on **LLDP Interface** and select **Create LLDP Interface Policy**.
3. Enter a descriptive **Name** for your LLDP interface policy.
4. Select **Enabled** for **Receive State**.
5. Select **Enabled** for **Transmit State**.
6. Click **Submit**.

STEP 2 | Create a Port Channel policy to enable LACP.

1. Select **Fabric** > **Access Policies** > **Interface Policies** > **Policies** > **Port Channel**.
2. Right-click on **Port Channel** and select **Create Port Channel Policy**.
3. Enter a descriptive **Name** for your port channel policy.
4. Select **LACP Active** from the **Mode** drop-down.
5. Click **Submit**.

Create an External Routed Network

The firewalls pass IP routing information to the ACI ovar a Layer 3 OSPF network. ACI uses a switch virtual interface (SVI) on the leaf switches with an IP address on each switch for connection resilience. Create a Layer 3 routed network to peer with the firewall using OSPF.

STEP 1 | On the **Tenants** tab, double-click on the name of your tenant.

STEP 2 | Select **Networking** > **External Routed Networks**.

STEP 3 | Right-click **External Routed Networks** and select **Create Routed Outside**.

STEP 4 | Enter a descriptive **Name** for your **External Routed Network**.

STEP 5 | Select your VRF with external connectivity from the **VRF** drop-down.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Select the external routed domain you created previously form the **External Routed Domain** drop-down.

STEP 7 | Select **OSPF**.

STEP 8 | Enter an **OSPF Area ID**. The Area ID can be expressed in decimal number or dotted decimal form. For example, Area 1 is the same as Area 0.0.0.1 or Area 271 is the same as Area 0.0.1.15. The Area ID range is 0 (0.0.0.0) to 4294967295 (255.255.255.255).

STEP 9 | Select **Regular Area** for the **OSPF Area Type**.

STEP 10 | Click the plus (+) button to the right of **Nodes and Interface Profiles** to create a Node Profile with a node that for the border-leaf switches that connect to the firewall.

STEP 11 | Enter a descriptive **Name** for your **Node Profile**.

STEP 12 | Attach nodes to your Node Profile.

1. Click the plus (+) button to the right of **Nodes**. This opens the **Select Node** window.
2. Select the node that your firewall is connected to from the **Node ID** drop-down.
3. Enter the IP address of the router attached to the leaf switch in **Router ID**.
4. Click **OK**.
5. Click the plus (+) button to the right of **Nodes and Interface Profiles**.
6. Enter a descriptive **Name** for your **Node Profile**.
7. Click the plus (+) button to the right of **Nodes**. This opens the **Select Node** window.
8. Select the node that your secondary HA firewall is connected to from the **Node ID** drop-down.
9. Enter the IP address of the router attached to the second leaf switch in **Router ID**.
10. Click **OK**.

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Attach an OSPF Interface Profile for your Node Profile.

1. Enter a descriptive **Name** for your OSPF Interface Profile.
2. Click **Next**.
3. Select **Create OSPF Interface Policy** from the OSPF Policy drop-down.
4. Enter a descriptive **Name** for your OSPF Interface Policy.
5. Select **MTU Ignore**.
6. Click **Submit**.
7. Click **Next**.
8. Click **SVI**.
9. Click the plus (+) button to the right of **SVI Interfaces**. This opens the **Select SVI** window.
10. Click **Virtual Port Channel**.
11. Select the Path to the port and port channel interface where the firewall connects to the leaf switch.
12. In **Encap**, enter the VLAN encapsulation used for your layer 3 outside profile.
13. Select **Trunk** for Mode.
14. In the **Side A IPv4 Primary Address** field, enter the primary IP address of the path attached to the layer 3 outside profile.
15. In the **Side B IPv4 Primary Address** field, enter the secondary IP address of the path attached to the layer 3 outside profile.
16. Click **OK**.

STEP 14 | Click **OK** to close the Create Interface Profile window.

STEP 15 | Click **OK** to close the Create Node Profile window.

STEP 16 | Click **Next**.

STEP 17 | Click the plus (+) button to the right of **External EPG Networks**. This opens the **Create Routed Outside** window.

STEP 18 | Enter a descriptive **Name** for you External Network.

STEP 19 | Add a subnet to you External Network.

1. Click the plus (+) button to the right of **Subnets**.
2. Enter the IP address and mask of the subnet's default gateway.
3. Select **Export Route Control Subnet**.
4. Select **External Subnets for External EPG**.
5. Click **OK**.

STEP 20 | Click **Finish**.

Configure Subnets to Advertise to the External Firewall

By default, subnets in the ACI fabric are not advertised to external networks. You must configure the subnets to be advertised externally.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | On the **Tenants** tab, double-click on the name of your tenant.

STEP 2 | Select **Networking** > **Bridge Domains** > **<your bridge domain>**.

STEP 3 | Click **L3 Configurations**.

STEP 4 | Click the plus (+) button to the right of **Associated L3 Outs**.

STEP 5 | Select the Layer 3 external routed network connection you created in the previous procedure from the **L3 Out** drop-down.

STEP 6 | Click **Update**.

STEP 7 | Select **Networking** > **Bridge Domains** > **<your bridge domain>** > **Subnets** > **<externally advertised subnet>**.

STEP 8 | Set the Scope to **Advertised Externally**.

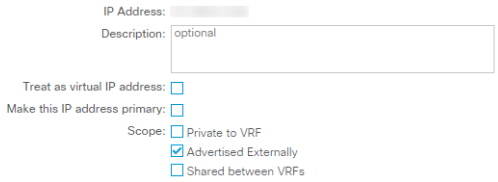

STEP 9 | Click **Submit**.

Create an Outbound Contract

Create a contract with a filter that allows DNS, NTP, HTTP, and HTTPS traffic. You will use this contract to allow all endpoints in the VRF to reach the external networks but limits the traffic sent to the firewall.

STEP 1 | On the **Tenants** tab, double-click on the name of your tenant.

STEP 2 | Select **Contracts** > **Filters**

STEP 3 | Right-click on **Filters** and select **Create Filter**.

STEP 4 | Enter a descriptive **Name** for the filter.

STEP 5 | Create a filter entry for UDP traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **UDP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **udp** from the **IP Protocol** drop-down.
5. Select **dns** from the **Destination Port From** drop-down.
6. Click **Update**.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Create a filter entry for TCP traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **TCP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **tcp** from the **IP Protocol** drop-down.
5. Select **dns** from the **Destination Port From** drop-down.
6. Click **Update**.

STEP 7 | Create a filter entry for NTP traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **NTP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **udp** from the **IP Protocol** drop-down.
5. In the **Destination Port From** field, enter 123.
6. Click **Update**.

STEP 8 | Create a filter entry for HTTP traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **HTTP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **tcp** from the **IP Protocol** drop-down.
5. Select **http** from the **Destination Port From** drop-down.
6. Click **Update**.

STEP 9 | Create a filter entry for HTTPS traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **HTTP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **tcp** from the **IP Protocol** drop-down.
5. Select **https** from the **Destination Port From** drop-down.
6. Click **Update**.

 ©2024 Palo Alto Networks, Inc.

STEP 10 | Click **Submit**.

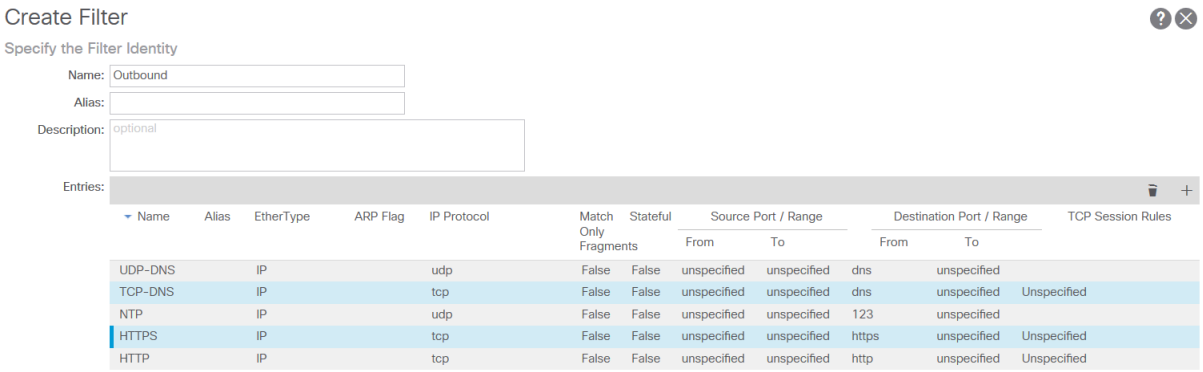

STEP 11 | Create a contract for outbound traffic.

1. On the **Tenants** tab, double-click on the name of your tenant and select **Contracts**.
2. Right-click on **Contracts** and select **Create Contract**.
3. Enter a descriptive **Name** for your **Contract**.
4. Click the plus (+) button to the right of **Subjects**.
5. Enter a descriptive **Name** for you **Subject**.
6. Under Filter Chain, click the plus (+) button to the right of **Filters**.
7. Select the filter you created previously from the drop-down.
8. Click **OK**.

STEP 12 | Click **Submit**.

Create an Inbound Web Contract

You must also create a contract and filters to allow inbound traffic to reach the servers behind the firewall. The following procedure describes the process of creating a contract and filters that allows HTTP and HTTPS web traffic to access resources behind the firewall.

STEP 1 | On the **Tenants** tab, double-click on the name of your tenant.

STEP 2 | Select **Contracts** > **Filters**

STEP 3 | Right-click on **Filters** and select **Create Filter**.

STEP 4 | Enter a descriptive **Name** for the filter.

STEP 5 | Create a filter entry for HTTP traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **HTTP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **tcp** from the **IP Protocol** drop-down.
5. Select **http** from the **Destination Port From** drop-down.
6. Click **Update**.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Create a filter entry for HTTPS traffic.

1. Click the plus (+) button to the right of **Entries**.
2. Enter a descriptive **Name** for the **TCP** filter.
3. Select **IP** from the **EtherType** drop-down.
4. Select **tcp** from the **IP Protocol** drop-down.
5. Select **https** from the **Destination Port From** drop-down.
6. Click **Update**.

STEP 7 | Click **Submit**.

STEP 8 | Create a contract for inbound web traffic.

1. On the **Tenants** tab, double-click on the name of your tenant and select **Contracts**.
2. Right-click on **Contracts** and select **Create Contract**.
3. Enter a descriptive **Name** for your **Contract**.
4. Click the plus (+) button to the right of **Subjects**.
5. Enter a descriptive **Name** for you **Subject**.
6. Under Filter Chain, click the plus (+) button to the right of **Filters**.
7. Select the filter you created previously from the drop-down.
8. Click **OK**.

STEP 9 | Click **Submit**.

Apply Outbound and Inbound Contracts to the EPGs

Now you must apply the inbound and outbound contracts to the appropriate EPGs.

For all the EPGs (EPG collection) within a VRF to send traffic to an external destination, each internal EPG must contract with the external EPG. Typically, you would need to create a separate contract between each internal EPG and the external EPG. However, using a vzAny object you can apply the same contract to all EPGs dynamically. The EPG collection consumes the contract and the external EPG provides the contract. You can configure specific traffic profiles in the contract or send all traffic to the firewall and allow it to control the traffic leaving the datacenter. Additionally, any new EPG that joins the VRF will automatically has the contract applied to it.

Apply the inbound contract so the internal EPG is the provider and the external EPG is the consumer. Traffic flowing to the internal EPG is fist checked against the contract and any allowed traffic is then secured further by the firewall as necessary.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Apply the outbound contract to all EPGs in the VRF.

1. On the **Tenants** tab, double-click on the name of your tenant.
2. Select **Networking** > **VRFs** > **<you VRF>** > **EPG Collection for VRF**.
3. Click the plus (+) button to the right of **Consumed Contracts**.
4. Select your outbound contract from the **Name** drop-down.
5. Click **Update**.
6. Select **Networking** > **External Routed Networks** > **<your external routed network>** > **Networks** > **External**.
7. Click the plus (+) button to the right of **Provided Contracts**.
8. Select your outbound contract from the **Name** drop-down.
9. Click **Update**.

STEP 2 | Apply the inbound contract so an internal EPG provides it to the external EPG.

1. On the **Tenants** tab, double-click on the name of your tenant.
2. Select **Application Profiles** > **<your application profile>** > **Application EPGs** > **<your application EPG>** > **Contracts**.
3. Right-click on **Contracts** and select **Add Provided Contract**.
4. Select your inbound contract from the **Contract** drop-down.
5. Click **Submit**.
6. On the same tenant, select **Networking** > **External Routed Networks** > **<your external routed network>** > **Networks** > **External**.
7. On the Contracts tab, click the plus (+) button to the right of **Consumed Contracts**.
8. Select your inbound contract from the **Name** drop-down.
9. Click **Update**.

Create a Virtual Router and Security Zone for North-South Traffic

Create a virtual router and security zone on the firewall to match the tenant and VRF on ACI.

STEP 1 | Log in to the firewall.

STEP 2 | Select **Network** > **Virtual Routers** and click **Add**.

STEP 3 | Give the virtual router a descriptive **Name**.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Click **OK**.

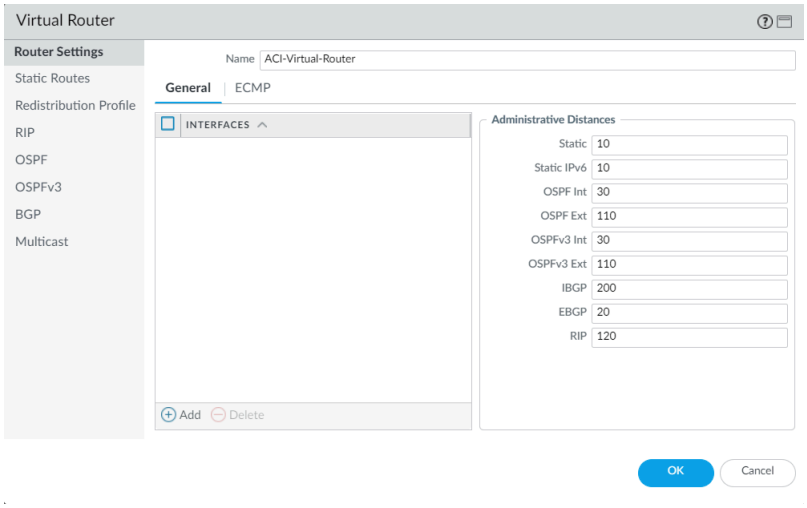

STEP 5 | Select **Network** > **Zones** and click **Add**.

STEP 6 | Give the zone a descriptive **Name**.

STEP 7 | Choose Layer 3 from the **Type** drop-down.

STEP 8 | Click **OK**.

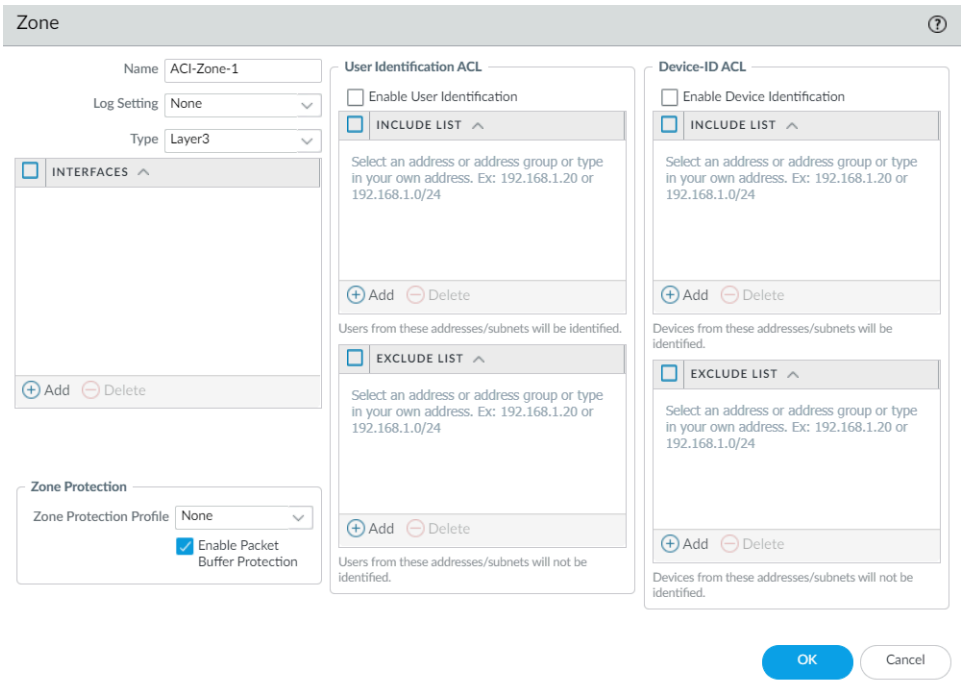

STEP 9 | **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

Configure the Network Interfaces

Configure an aggregate Ethernet interface, member interfaces, and subinterface that your firewall uses to connect to the ACI leaf switches. If you are using a VM-Series firewall, use discreet interfaces instead of aggregate interfaces.

The VM-Series firewall does not support aggregate Ethernet groups.

STEP 1 | Select **Network** > **Interfaces** > **Ethernet** and click **Add Aggregate Group**.

STEP 2 | Enter a number for the aggregate group in the second **Interface Name** field.

STEP 3 | Select Layer 3 from the **Interface Type** drop-down.

STEP 4 | Select the **LACP** tab and click **Enable LACP**.

STEP 5 | Select **Fast** as the **Transmission Rate**.

STEP 6 | Under High Availability Options, select **Enable in HA Passive State**.

*Do not select **Same System MAC Address for Active-Passive HA**. This option makes the firewall pair appear as a single device to the switch, so traffic will flow to both firewalls instead of just the active firewall.*

STEP 7 | Click **OK**.

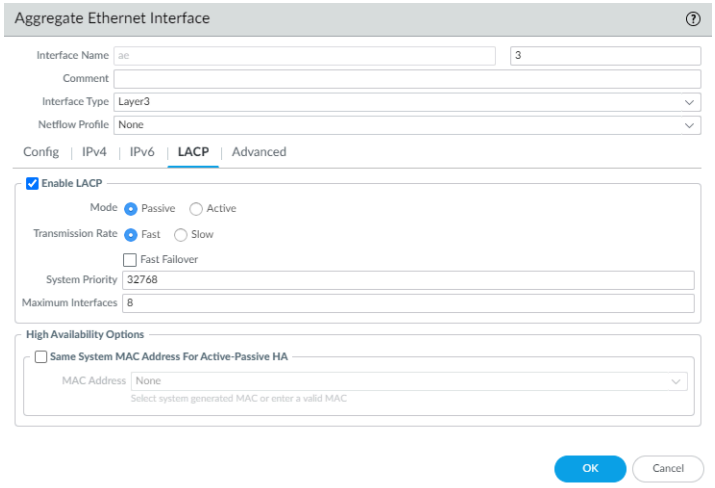

 ©2024 Palo Alto Networks, Inc.

STEP 8 | Click on the name of an Ethernet interface to configure it and add it to the aggregate group.

1. Select **Aggregate Ethernet** from the Interface Type drop-down.
2. Select the interface you defined in the aggregate Ethernet group configuration.
3. Click **OK**.
4. Repeat this step for each other member interface of the aggregate Ethernet group.

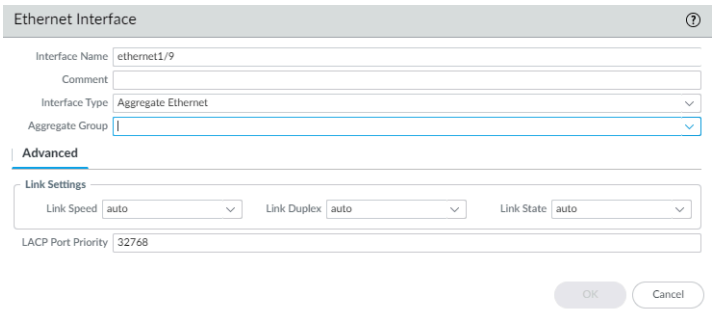

STEP 9 | Add a subinterface on the aggregate Ethernet interface for the tenant and VRF.

1. Select the row of your aggregate Ethernet group and click **Add Subinterface**.
2. In the second **Interface Name** field, enter a numerical suffix to identify the subinterface.
3. In the **Tag** field, enter the VLAN tag of the subinterface.
4. Select the virtual router you configured previously from the **Virtual Router** drop-down.
5. Select the zone you configured previously from the **Zone** drop-down.
6. Select the **IPv4** tab.
7. Select the **Static** Type.
8. Click **Add** and enter the subinterface IP address and network mask in CIDR notation.
9. Click **OK**.

Configure Route Redistribution and OSPF

Configure route redistribution to make routing information from the firewall available to the external-facing routers attached to your leaf switches. Then configure OSPF on the firewall and assign a router-id, area number, and interface to form adjacencies.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Configure route redistribution.

1. Select **Network** > **Virtual Routers** and click on the virtual router you created earlier.
2. Select **Redistribution Profile** > **IPv4** > **Add**.
3. Enter a descriptive **Name** for your redistribution profile.
4. Enter a priority.
5. For **Redistribute**, select **Redist**.
6. Check **connect** and **static** under **General Filters**.
7. Click **OK**.

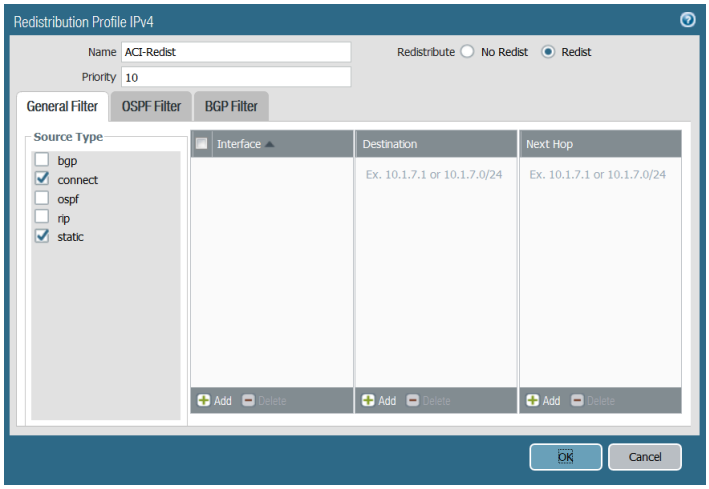

STEP 2 | Configure OSPF.

1. Select **Network** > **Virtual Routers** and click on the virtual router you created earlier.
2. Select **Router Settings** > **ECMP** and select **Enable**.
3. Select **OSPF** and choose **Enable**.
4. Enter the OSPF **Router ID**.
5. Under **Area**, click **Add**.
6. Enter the **Area ID**. This value must match the value you assigned when you created the external routed network on the APIC. On the firewall, this must be entered in dotted

 ©2024 Palo Alto Networks, Inc.

decimal form. For example, if you entered an Area ID of 10 in the APIC, the equivalent on the firewall is 0.0.0.10.

7. Select **Interface** > **Add**.
8. Enter the interface that connects to your external network EPG and click **OK**.
9. Select **Export Rules** > **Add**.
10. Select the Redistribution Profile you created above from the **Name** drop-down and click **OK**.
11. Select **Allow Redistribute Default Route**.
12. Click **OK**.

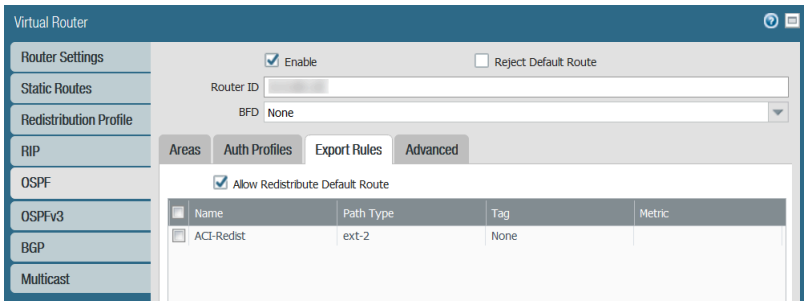

Configure NAT for External Connections

You only need to configure NAT if the firewall has an external interface used for connecting to networks outside of your data center. While NAT is not required, you can use this procedure to translate private IP addressing in your data center to public IP addressing outside. Begin setting up NAT by configuring address translation for traffic entering server inside an EPG in your data center. Then configure a NAT policy that translates the source address of outbound traffic from any EPG to the external interface IP address.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Configure address translation for traffic entering an EPG in your data center.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a descriptive **Name** for your NAT policy rule.
3. Select **Original Packet** and click **Add** under **Source Zone**.
4. Select the source zone from the drop-down.
5. Select the destination zone from the **Destination Zone** drop-down.
6. Select **Any** for the **Source Address**.
7. Click **Add** under **Destination Address** and enter the external IP address.

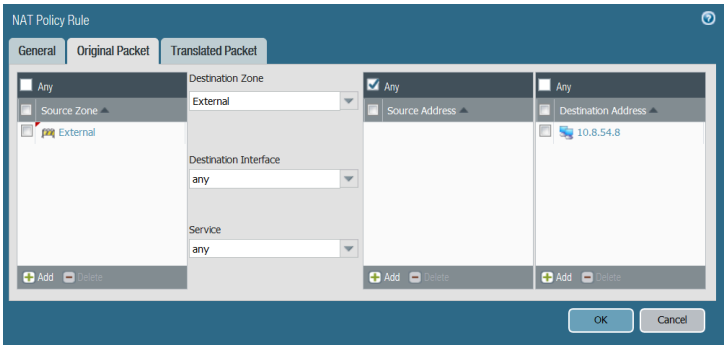

8. On the **Translated Packet** tab, select the **Translation Type** under **Destination Address Translation**.
9. Select an address from the **Translated Address** drop-down.
10. Click **OK**.

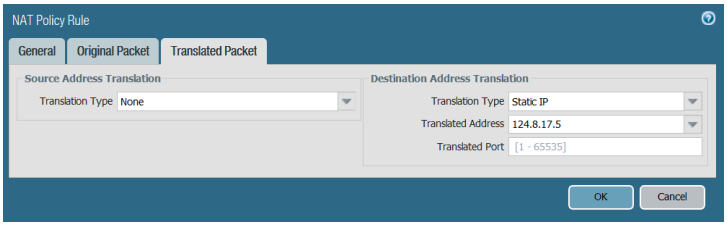

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Configure address translation for outbound traffic.

1. Select **Policies** > **NAT** and click **Add**.
2. Enter a descriptive **Name** for your outbound NAT policy.
3. Select **Original Packet** and click **Add** under **Source Zone**.
4. Select the zone that matches your ACI tenant and VRF.
5. Select the external zone from the **Destination Zone** drop-down.

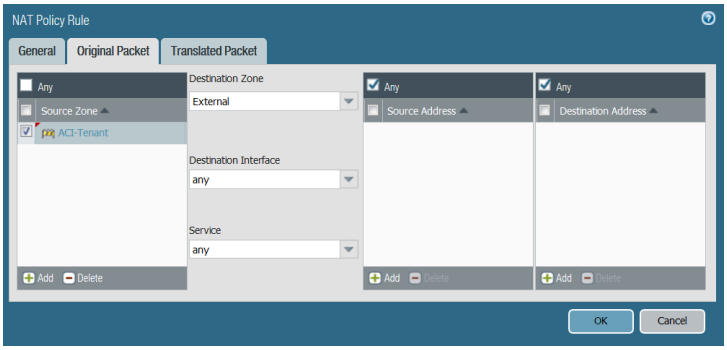

6. On the **Translated Packet** tab, select the **Translation Type** under **Source Address Translation**.
7. Enter additional required address information.
8. Click **OK**.

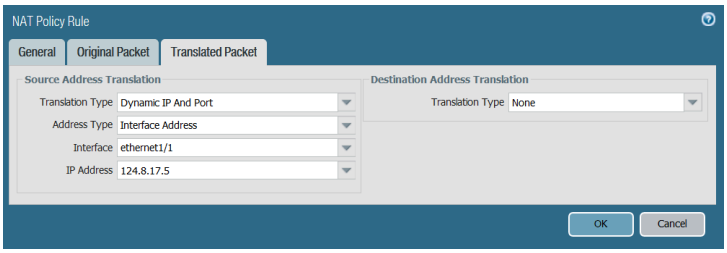

STEP 3 | **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

Endpoint Monitoring in Cisco ACI

The Cisco ACI plugin for Panorama allows you to build security policy for your Cisco ACI fabric using Dynamic Address Groups. The plugin monitors for changes in an Application Policy Infrastructure Controller (APIC) fabric in your Cisco ACI environment and shares that information with Panorama. Each Panorama with the Cisco ACI plugin installed can support up to 16 APIC clusters. And each monitoring definition has one cluster and one notify group.

The number of endpoints that the Cisco ACI plugin can monitor is dependent the amount of memory allocated to Panorama. If you have a Panorama virtual appliance, make sure you assign the necessary amount of memory for the endpoints in your environment. See the Panorama Admin Guide for more information about preparing your virtual Panorama.

Panorama Memory	Endpoints
8GB	10,000
16GB	20,000

The Cisco ACI plugin processes the endpoint information and converts it into a set of tags that can be used as match criteria for placing IP addresses in dynamic address groups. The tags are constructed in the following format:

cisco.cl_<cluster>.tn_<tenant>.ap_<app-profile>.{epg_<EPG> | uepg_<micro-EPG>}

- **cisco.cl_<cluster>**—this tag groups IP addresses into a dynamic address group based on the Cisco ACI cluster and displays the name of your cluster.
- **cisco.cl_<cluster>.tn_<tenant>**—this tag groups IP addresses into a dynamic address group based on tenant and displays the name of your cluster and tenant.
- **cisco.cl_<cluster>.tn_<tenant>.ap_<app-profile>**—this tag groups IP addresses into a dynamic address group base on application profile and displays the name of your cluster, tenant, and application profile.
- **cisco.cl_<cluster>.tn_<tenant>.ap_<app-profile>.epg_<EPG>**—this tag groups IP addresses into a dynamic address group based on EPG and displays the name of your cluster, tenant, application profile, and EPG.
- **cisco.cl_<cluster>.tn_<tenant>.ap_<app-profile>.uepg_<micro-EPG>**—this tag groups IP addresses into a dynamic address group based on micro-EPG and displays the name of your cluster, tenant, application profile, and micro-EPG.
- **cisco.cl_<cluster>.tn_<tenant>.l2out_<L2-external-endpoint>**—this tag groups IP addresses into dynamic address groups based on L2 external endpoint and displays the name of you cluster, tenant, and L2 external endpoint.
- **cisco.cl_<cluster>.tn_<tenant>.bd_<bridge-domain>.subnet_<subnet>**—this tag groups IP address into a dynamic address group based on subnet and displays the name of you cluster, tenant, bridge domain, and subnet.

To retrieve endpoint IP-address-to-tag mapping information, you must configure a Monitoring Definition for each APIC fabric in your Cisco ACI environment. The Monitoring Definition specifies the username and password that allows Panorama to connect to the APICs. It also

 ©2024 Palo Alto Networks, Inc.

specifies the device groups and corresponding notify groups containing the firewalls to which Panorama pushes the tags. After you configure the Monitoring Definition and the Cisco ACI plugin retrieves the tags, you can create dynamic address groups and add the tags as match criteria.

The Cisco ACI plugin uses two intervals to retrieve information from the APIC. The first is the monitoring interval.

- Monitoring interval—The monitoring interval is the amount of time that the plugin waits before querying for changes in the fabric. If no changes occurred, the monitoring interval resets. If changes are detected, the plugin processes the changes before resetting the monitoring interval. The default monitoring interval is 60 seconds. You can set the monitoring interval from 60 seconds to one day (86,400 seconds).
- Full-sync interval—The full-sync interval is the amount of time that the plugin waits before updating the dynamic objects from all fabrics regardless of any changes occurred. This ensures that the plugin is synchronized with the fabric even if a change event is missed by the monitoring interval. The default full-sync interval is 10 minutes. You can set the full-sync interval from 600 seconds (10 minutes) and 86,400 seconds (one day).

  You must configure the full-sync interval through the Panorama CLI.

If you configure a value for the monitoring interval greater than that of the full-sync interval, the full-sync interval is ignored and a full synchronization is performed at every monitoring interval.

If Panorama loses its connection with the APIC, Panorama will attempt to reconnect five times. After five failed attempts, Panorama stops monitoring for changes in your clusters and displays the reconnection attempts in the system log. To recover and begin monitoring your clusters again, you must perform a commit on Panorama.

- Install the Panorama Plugin for Cisco ACI
- Configure the Cisco ACI Plugin
- Panorama Plugin for Cisco ACI Dashboard

Install the Panorama Plugin for Cisco ACI

To get started with endpoint monitoring on Cisco ACI, download and install the Cisco ACI plugin on Panorama.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panoramas in an HA pair, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

 ©2024 Palo Alto Networks, Inc.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Verify that your virtual Panorama has enough memory to support the number endpoints in your ACI environment.

STEP 2 | Select **Panorama** > **Plugins**.

STEP 3 | Select **Check Now** to retrieve a list of available updates.

STEP 4 | Select **Download** in the Action column to download the plugin.

STEP 5 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

Configure the Cisco ACI Plugin

After installing the plugin, you must set the monitoring interval, configure a notify group, and establish a connection between Panorama and the APIC fabric.

STEP 1 | (Optional) Configure the full-sync interval.

1. Log in to the Panorama CLI.
2. Enter configure mode.

   `admin@Panorama> configure`
3. Use the following command to set the full-sync interval. The default interval is 600 seconds (10 minutes). The range is 600 seconds to 86,400 seconds (one day).

   admin@Panorama# **`set plugins cisco full-sync-interval <interval-in-seconds>`**

STEP 2 | Log in to the Panorama web interface.

STEP 3 | You must add the firewalls as managed devices on Panorama and create Device Groups so that you can configure Panorama to notify these groups with the VM information it retrieves. Device groups can include VM-Series firewalls or virtual systems on the hardware firewalls.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Enable monitoring, set the monitoring interval, and enable bypass proxy.

1. Select **Panorama** > **Cisco ACI** > **Setup** > **General**.
2. Select **Enable Monitoring**. This enables monitoring for all clusters in your deployment.
3. Set the **Monitoring Interval** in seconds. The monitoring interval is how often Panorama retrieves updated network information from the APIC. The default value is 60 seconds and the range is 60 seconds to 86,400 seconds (one day).
4. (Optional) Select Bypass Proxy to Bypass proxy server settings, configured on Panorama under **Panorama** > **Setup** > **Services** > **Proxy Server**, for communication between Panorama and the APIC. This allows Panorama to communicate directly with the APIC while maintaining proxied communication for other services.

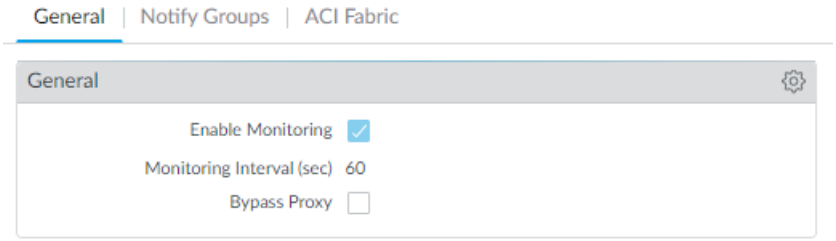

STEP 5 | Create a notify group.

1. Select **Panorama** > **Cisco ACI** > **Setup** > **Notify Groups**.
2. Click **Add**.
3. Enter a descriptive **Name** for your notify group.
4. Select the device groups in your ACI deployment.

 ©2024 Palo Alto Networks, Inc.

STEP 6 | Add ACI fabric information.

1. Select **Panorama** > **Cisco ACI** > **Setup** > **ACI Fabric**.
2. Enter a descriptive **Name** for your cluster.
3. Enter the IP address or FQDN for each APIC in the cluster as a comma-separated list.

   When using FQDN, do not include https:// in the URL.
4. Enter your APIC username.
5. Enter and confirm your APIC password.
6. Click **OK**.

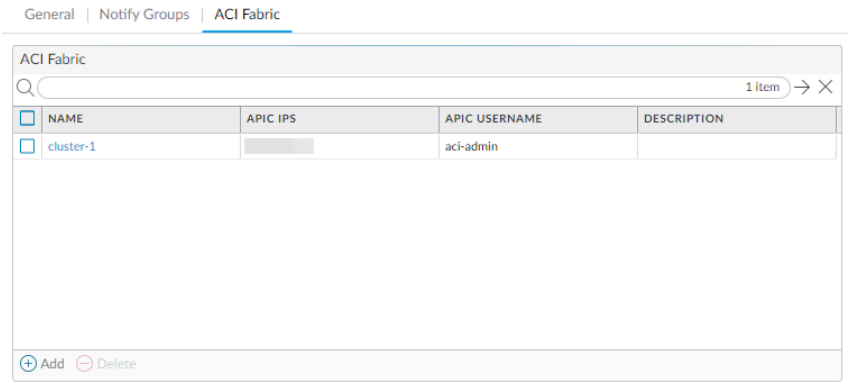

STEP 7 | Configure the Monitoring Definition.

1. Select **Panorama** > **Cisco ACI** > **Monitoring Definition** and click **Add**.
2. Enter a descriptive **Name** and optionally a description to identify the Cisco ACI cluster for which you use this definition.
3. Select the **Cluster Info** and **Notify Group**.
4. Click **OK**.

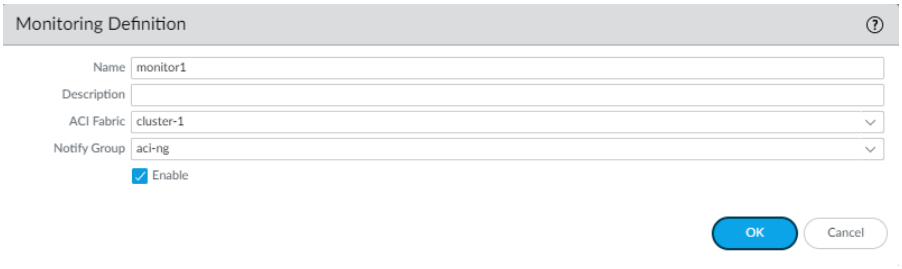

STEP 8 | **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Verify that you can view the EPG information on Panorama, and define the match criteria for Dynamic Address Groups.

Some browser extensions may block API calls between Panorama and the APIC which prevents Panorama from receiving match criteria. If Panorama displays no match criteria and you are using browser extensions, disable the extensions and Synchronize Dynamic Objects to populate the tags available to Panorama.

Panorama does not immediately process new monitoring definitions and populate the match criteria available to dynamic address. You should wait for the duration of your configured monitoring interval before verifying that EPG information.

STEP 10 | Verify that addresses in your EPGs are added to dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click **More** in the Addresses column of a dynamic address group.

   Panorama displays a list of IP addresses added to that dynamic address group based on the match criteria you specified.

STEP 11 | Use dynamic address groups in policy.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

   See Use Dynamic Address Groups in Policy for more information.

STEP 12 | You can update the dynamic objects from the APIC at any time by synchronizing dynamic objects. Synchronizing dynamic objects enables you to maintain context on changes in the

 ©2024 Palo Alto Networks, Inc.

virtual environment and allows you to enable applications by automatically updating the Dynamic Address Groups used in policy rules.

1. Select **Panorama** > **Cisco ACI** > **Monitoring Definition**.
2. Click **Synchronize Dynamic Objects**.

On HA failover, the newly active Panorama attempts to reconnect to the APIC and retrieve tags for all monitoring definitions. If there is an error with reconnecting even one monitoring definition, Panorama generates a system log message

```
Unable to process subscriptions after HA switch-over;
 user-intervention required.
```

When you see this error, you must log in to Panorama and fix the issue, for example remove an invalid APIC IP or provide valid credentials, and commit your changes to enable Panorama to reconnect and retrieve the tags for all monitoring definitions. Even when Panorama is disconnected from the APIC, the firewalls have the list of all tags that had been retrieved before failover, and can continue to enforce policy on that list of IP addresses. If you perform a commit before resolving the failover error, the newly active Panorama will not push any IP-to-tag mapping information and clearing the mapping information from the firewalls. As a best practice, to monitor this issue, configure action-oriented log forwarding to an HTTPS destination *from Panorama so that you can take immediate action.*

Panorama Plugin for Cisco ACI Dashboard

The Panorama plugin for Cisco ACI dashboard provides a bird's-eye view of the ACI infrastructure monitored by the plugin. The dashboard consists of two pages—the first page provides an overview of various objects monitored by the plugin on a set of clickable tiles; clicking a tile takes you to the second page that provides further information about the object displayed on the tile.

After installing the plugin, you can access the dashboard by selecting **Panorama** > **Cisco ACI** > **Dashboard**.

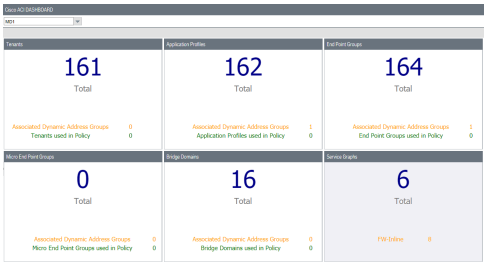

The dashboard only queries for and counts pre-rule security policies configured on Panorama; it does not include post-rules, default-rules, or NAT rules.

Dashboard Tiles	Description
Tenant Tags	Displays the total number of tenants Panorama retrieved from the APIC. Additionally, it displays the number of dynamic address groups associated with tenants and the number of tenants used in policy.

 ©2024 Palo Alto Networks, Inc.

Dashboard Tiles	Description
If a tenant's health score is zero (0) on the APIC, Panorama does not retrieve information about that tenant. Therefore, the tenant count on Panorama will not match the total on the APIC.	Click the tile to drill down and view the following columns. • **Tenant Name**—list all the tenants retrieved by Panorama. • **Tenant Tag**—the Panorama tag associated with each tenant. • **Dynamic Address Group**—displays the dynamic address groups associated with the listed tag. • **In Policy**—shows if the listed dynamic address group is used in policy.
Application Profiles	Displays the total number of application profiles Panorama retrieved from the APIC. Additionally, it displays the number of dynamic address groups associated with application profiles and the number of application profiles used in policy. Click the tile to drill down and view the following columns. • **Application Profile Name**—lists all application profiles retrieved by Panorama. • **Tenant Name**—displays the tenant associated with the listed application profile. • **Application Profile Tag**—the Panorama tag associated with each application profile. • **Dynamic Address Group**—displays the dynamic address groups associated with the listed tag. • **In Policy**—shows if the listed dynamic address group is used in policy.

 ©2024 Palo Alto Networks, Inc.

Dashboard Tiles	Description
End Point Groups	Displays the total number of end point groups (EPG) Panorama retrieved from the APIC. Additionally, it displays the number of dynamic address groups associated with EPGs and the number of EPGs used in policy. Click the tile to drill down and view the following columns. • **EPG Name**—lists all EPGs retrieved by Panorama. • **Application Profile Name**—lists the EPG's associated application profile. • **Tenant Name**—displays the tenant associated with the listed application profile. • **EPG Tag**—the Panorama tag associated with each EPG. • **Dynamic Address Group**—displays the dynamic address groups associated with the listed tag. • **In Policy**—shows if the listed dynamic address group is used in policy.
Micro End Point Groups	Displays the total number of micro end point groups (EPG) Panorama retrieved from the APIC. Additionally, it displays the number of dynamic address groups associated with micro EPGs and the number of micro EPGs used in policy. Click the tile to drill down and view the following columns. • **Micro EPG Name**—lists all EPGs retrieved by Panorama. • **Application Profile Name**—lists the Micro EPG's associated application profile. • **Tenant Tag**—displays the tenant associated with the listed application profile. • **Micro EPG Tag**—the Panorama tag associated with each Micro EPG. • **Dynamic Address Group**—displays the dynamic address groups associated with the listed tag. • **In Policy**—shows if the listed dynamic address group is used in policy.
Bridge Domains	Displays the total number of bridge domains Panorama retrieved from the APIC. Additionally, it displays the number of dynamic address groups associated with brdige domains and the number of bridge domains used in policy. Click the tile to drill down and view the following columns. • **Bridge Domain Name**—lists all bridge domains retrieved by Panorama. • **Tenant Name**—displays the tenant associated with the listed bridge domain.

 ©2024 Palo Alto Networks, Inc.

Dashboard Tiles	Description
	• **Bridge Domain Tag**—the Panorama tag associated with each bridge domain. • **Dynamic Address Group**—displays the dynamic address groups associated with the listed tag. • **In Policy**—shows if the listed dynamic address group is used in policy.
Service Graphs	Displays the total number of Service Graphs monitored by the plugin as well as well as the number of firewalls in line with monitored service graphs. Click the tile to drill down and view the following columns. • **Service Graph Name**—lists all service graphs retrieved by Panorama. • **Producer EPG**—displays the producer EPG associated with the service graph. • **FW InLine**—displays the firewall associated with the service graph.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on Cisco CSP

You can deploy the VM-Series firewall as a network virtual service on the Cisco Cloud Security Platform (CSP). Because the Cisco CSP is a RHEL KVM platform, the VM-Series firewall is deployed using the VM-Series firewall for KVM base image.

With the VM-Series firewall on Cisco CSP, you can protect your workloads, prevent advanced threats, and improve visibility into the applications on your virtual network.

- VM-Series on Cisco CSP System Requirements
- Deploy the VM-Series Firewall on Cisco CSP

VM-Series on Cisco CSP System Requirements

You can create and deploy multiple instances—standalone or as an HA pair—of the VM-Series firewall on your Cisco CSP.

The VM-Series firewall has the following requirements:

- See the Compatibility Matrix for supported versions of CSP and PAN-OS.
- Bootstrap Package converted to a ISO file
- See VM-Series System Requirements for the minimum hardware requirements for your VM-Series model.
- Minimum of two network interfaces (vNICs). One is a dedicated vNIC for the management interface and one is for the data interface. You can then add up to eight more vNICs for data traffic.
- The VM-Series firewall on Cisco CSP supports all VM-Series models except the VM-50.
- SR-IOV, packet MMAP, and DPDK mode are supported.

 ©2024 Palo Alto Networks, Inc.

Deploy the VM-Series Firewall on Cisco CSP

Complete the following procedure to deploy the VM-Series firewall on Cisco CSP.

STEP 1 | Download the VM-Series qcow2 base image file from the Customer Support Portal.

STEP 2 | Create a bootstrap ISO file for the VM-Series firewall.

1. Create a Bootstrap Package for your VM-Series firewall.
2. Create an ISO file containing the bootstrap package using your preferred tool.

STEP 3 | Log in to the Cisco CSP web interface.

STEP 4 | Upload the VM-Series firewall qcow2 image and ISO file.

1. Select **Configuration** > **Repository**.
2. Click the plus (+) icon.
3. Click **Browse** and navigate to your qcow2 file.
4. Click **Upload**.
5. Click **Browse** and navigate to your ISO file.
6. Click **Upload**.

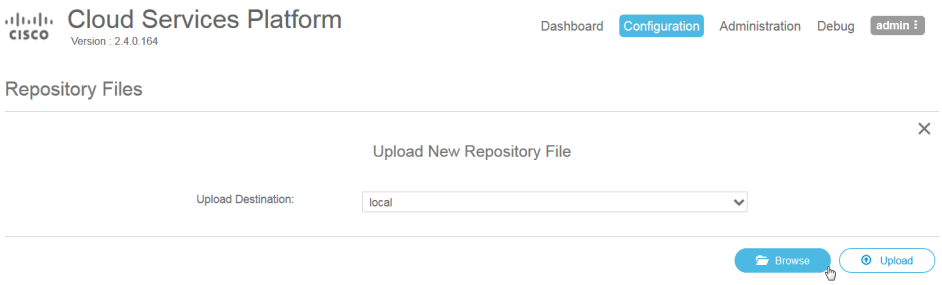

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Create the VM-Series firewall service.

1. Enter a descriptive **Name** for the VM-Series firewall.
2. Select the **Target Host Name** from the drop-down.
3. Select the qcow2 file you uploaded from the **Image Name** drop-down.

4. Select the Day Zero Config.
   1. Click the **Day Zero Config** plus (+) icon.
   2. Select the bootstrap ISO file from the **Source File Name** drop-down.
   3. Click **Submit**.

5. Allocate the number of cores and memory required for your VM-Series firewall model.
6. Add enough vNICs to support the number of VM-Series interfaces configured in your bootstrap ISO file.

See the Cisco Cloud Service Platform documentation for more information about creating and deploying a service instance.

STEP 6 | After the bootstrap process is complete, log in to your VM-Series firewall using the management IP address your specified in the bootstrap ISO file.

The firewall should be up and configure based on the parameters you defined in the bootstrap package.

 ©2024 Palo Alto Networks, Inc.

Endpoint Monitoring for Cisco TrustSec

Install and configure the Panorama plugin for Cisco TrustSec to retrieve the IP addresses of endpoints in your environment and build security policy for those endpoints using Dynamic Address Groups.

- Panorama Plugin for Cisco TrustSec
- Install the Panorama Plugin for Cisco TrustSec
- Configure the Panorama Plugin for Cisco TrustSec
- Troubleshoot the Panorama Plugin for Cisco TrustSec

Panorama Plugin for Cisco TrustSec

The Panorama plugin for Cisco TrustSec enables you to create security policy for your TrustSec environment using dynamic address groups. The plugin monitors for changes in TrustSec security groups and registers that information with Panorama and forwards IP information to the firewall, so Panorama can apply the correct policy to corresponding endpoints. The Panorama plugin for Cisco TrustSec supports up to 16 pxGrid (Cisco ISE) servers.

The Panorama plugin for Cisco TrustSec currently supports dynamic objects but not static objects.

The Panorama plugin processes the endpoint information and converts it to a set of tags that you can use as match criteria for placing IP addresses in dynamic address groups. Panorama creates a tag for each security group tag (SGT) on your pxGrid servers. The tags are constructed in the following format:

cts.svr_<pxgrid-server-name>.sgt_<SGT-name>

To retrieve endpoint IP-address-to-tag mapping information, you must configure a Monitoring Definition for each pxGrid server in your environment. The pxGrid server configuration specifies the username and password and is referenced by the monitoring definition that allows Panorama to connect to the pxGrid. Additionally, you can configure the plugin to verify the pxGrid server identity with a certificate profile on Panorama. It also specifies the device groups and corresponding notify groups containing the firewalls to which Panorama pushes the tags. After you configure the Monitoring Definition and the plugin retrieves the tags, you can create dynamic address groups and add the tags as match criteria.

The Panorama Plugin for Cisco TrustSec version 1.0.2 and later supports Bulk Sync and PubSub monitoring modes. The plugin selects a mode based on the Panorama version—Bulk Sync mode if the Panorama version is earlier than 10.0.0, and PubSub mode on Panorama 10.0.0 and later. The user interface displays the configuration options for the default monitoring mode.

- Bulk Sync
- PubSub

Bulk Sync

Bulk Sync mode uses two intervals to retrieve information from your pxGrid servers—the monitoring interval and full-sync interval. This mode is the default when the Panorama Plugin for Cisco TrustSec version 1.0.2 or later is installed on a Panorama version earlier than 10.0.0. Panorama versions earlier than 10.0.0 support IP-tab updates to configd every 10 seconds.

- **Monitoring interval**—The monitoring interval is the amount of time that the plugin waits before querying for changes. If no changes have occurred, the monitoring interval resets. If there are changes, the plugin processes the changes before resetting the monitoring interval. The default monitoring interval is 60 seconds. You can set the monitoring interval from 10 seconds to one day (86,400 seconds).

  The minimum monitoring interval is 30 seconds when the Panorama plugin for Cisco TrustSec 1.0.0 is installed.

 ©2024 Palo Alto Networks, Inc.

- **Full-sync interval**—The full-sync interval is the amount of time that the plugin waits before updating the dynamic objects from all pxGrid servers regardless of any changes occurred. This ensures that the plugin is synchronized with the pxGrid server even if a change event is missed by the monitoring interval. You can set the full-sync interval from 600 seconds (10 minutes) to 86,400 seconds (one day). You must configure the full-sync interval from the Panorama CLI.

If the monitoring interval is greater than the full-sync interval, the full-sync interval is ignored and a full synchronization is performed at every monitoring interval.

PubSub

PubSub mode monitors notifications directly from the Cisco ISE server (the subscription daemon), parses for IP tags, and sends relevant information to the tag processing daemon (tag-proc). PubSub is the default mode when the Panorama Plugin for Cisco TrustSec version 1.0.2 or later is installed on Panorama version 10.0.0 or later. Panorama versions 10.0.0 or later support IP-tab updates to configd every 100 milliseconds.

- **Push interval**—The push interval is the amount of time between pushes. If the previous push takes too much time, the next push is triggered as soon as it finishes. The minimum push interval is 100 milliseconds (0 seconds) and the maximum is 60 seconds. The default push interval is 0 seconds.
- **Enable Full Sync**—Enable this option to trigger a complete update. If you enable full sync, you can set the full-sync interval. Default is no.
- **Full-sync interval**—The full-sync interval is the amount of time that the plugin waits before updating the dynamic objects from all pxGrid servers regardless of any changes occurred. The default full-sync interval is 10 minutes. You can set the full-sync interval from 600 seconds (10 minutes) to 86,400 seconds (one day). You must configure the full-sync interval from the Panorama CLI.
- **Reconnection interval**—The initial reconnection interval is 1 second, and it is doubled if the previous reconnection failed. The maximum reconnection interval is 64 sec. There is no limit to the number of reconnection attempts.

 ©2024 Palo Alto Networks, Inc.

Install the Panorama Plugin for Cisco TrustSec

To get started with endpoint monitoring with Cisco TrustSec, download and install the Cisco TrustSec plugin on Panorama. To correlate the plugin version with the Panorama version, see Panorama Plugins in the Compatibility Matrix.

Cisco TrustSec plugin upgrade or downgrade requires a commit.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panorama appliances in an HA pair, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Select **Panorama** > **Plugins**.

STEP 2 | Click **Check Now** to get the latest version of the plugin.

STEP 3 | Select **Download** in the Action column to download the plugin.

STEP 4 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

 ©2024 Palo Alto Networks, Inc.

Configure the Panorama Plugin for Cisco TrustSec

After you install the plugin, you must also assign a notify group to your Cisco TrustSec plugin configuration. A notify group is a list of device groups that includes the firewalls to which Panorama should push all the tags it retrieves from the pxGrid server.

Each Panorama with the Cisco TrustSec plugin installed can support up to 16 pxGrid servers and 16 monitoring definitions. And each monitoring definition has one pxGrid server and one notify group.

The Panorama plugin for Cisco TrustSec currently supports dynamic objects but not static objects.

The following configuration instructions cover both Bulk Sync and PubSub monitoring modes; some user interface features are enabled, or visible based on the monitoring mode.

STEP 1 | Configure the full-sync interval if you want to change it from the default 600 seconds (10 minutes).

1. Log in to the Panorama CLI.
2. Enter configure mode.

   ```
   admin@Panorama> configure
   ```
3. Use the following command to set the full-sync interval. The range is 600 seconds to 86,400 seconds (one day).

   ```
   admin@Panorama# set plugins cisco_trustsec full-sync-interval
   <interval-in-seconds>
   ```

STEP 2 | Log in to the Panorama web interface.

STEP 3 | You must add the firewalls as managed devices on Panorama and create Device Groups so that you can configure Panorama to notify these groups with the VM information it retrieves. Device groups can include VM-Series firewalls or virtual systems on the hardware firewalls.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Configure Cisco TrustSec monitoring.

1. Select **Panorama** > **Cisco TrustSec** > **Setup** > **General**.

   Enable Cisco TrustSec Monitoring is enabled by default. This enables monitoring for all clusters in your deployment.

   The user interface selects the PubSub monitoring mode if the Panorama plugin for Cisco TrustSec is 1.0.2 or later is installed on Panorama 10.0.0 or later:

   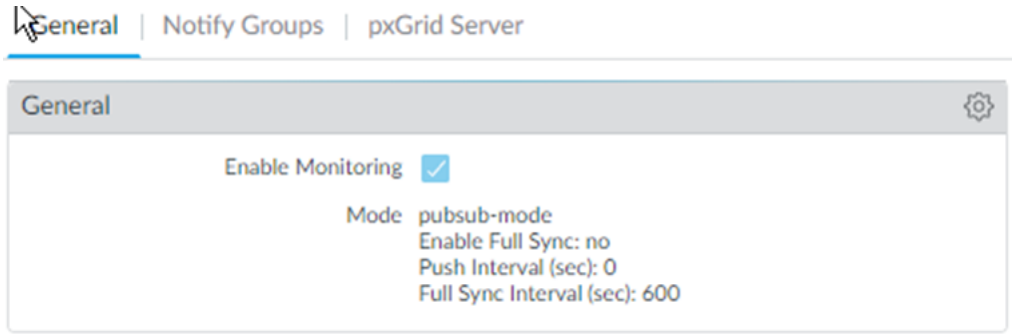

   The plugin selects Bulk Sync mode when it is installed on a Panorama version earlier than 10.0.0:

   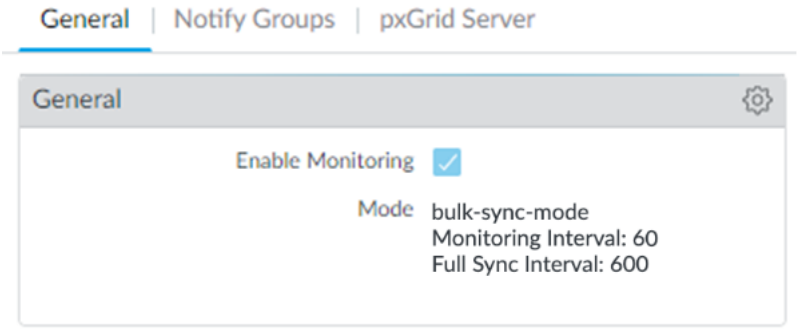

2. Click the gear to edit the setup parameters.

   - **Push Interval** (PubSub only)—Minimum 0, maximum 60 seconds, default is 0 (100 milliseconds).
   - **Enable Full Sync** (PubSub only, optional)—Select this optional to enable a full sync. Default is no.
   - **Full Sync Interval**.
     - PubSub—If Enable Full Sync is selected, you can set the full sync interval in seconds. Range is 600 seconds to 86400 seconds (one day), and the default value is 600
     - Bulk Sync—Enabled by default in Bulk Sync mode. Range is 600 seconds to 86400 seconds (one day), and the default value is 600.
   - **Monitoring Interval** (Bulk Sync only)—10 to 86400 seconds, default is 60—Set the polling interval at which Panorama queries the pxGrid server for endpoint address information. This is the time period between the end of a monitoring event and the start of the next event.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | Create a notify group.

1. Select **Panorama** > **Cisco TrustSec** > **Setup** > **Notify Groups**.
2. Click **Add**.
3. Enter a descriptive **Name** for your notify group.
4. Select the device groups you created in previously.

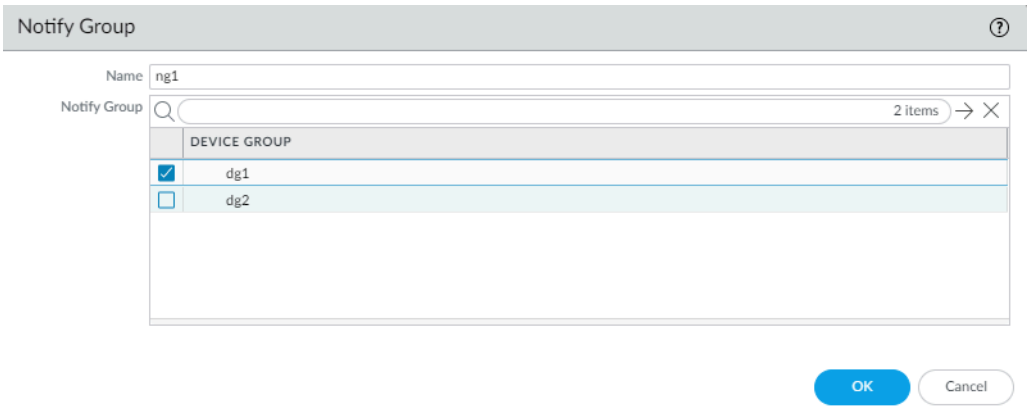

STEP 6 | (Optional) If enabling server identity verification of the pxGrid server, configure a certificate profile on Panorama.

STEP 7 | Create, activate, and approve the pxGrid client name and client password.

1. Log in to the Panorama CLI.
2. Execute the following command to create the client name.
   - If you have a certificate profile, create the client name as follows:

     admin@Panorama> request plugins cisco_trustsec create-account client-name <client-name> host <ise-server-ip>
   - If you skipped Step 6 and you do not have a certificate, enter:

     request plugins cisco_trustsec create-account server-cert-verification-enabled no client-name <client-name>host <host-name>
3. Execute the following command to create the client name.

```
admin@Panorama> request plugins cisco_trustsec create-account
 client-name test host 10.10.10.15

AccountCreate in progress...
AccountCreate successful.
  client nodename: test
  client password: PmVKBmPgf63Hypq

AccountActivate in progress...
AccountActivate successful.
```

 ©2024 Palo Alto Networks, Inc.

```
Please approve the account on the server.
```

4. Log in to your Cisco ISE server to approve the account.
5. Select **Administration** > **pxGrid Services** > **All Clients**.
6. Select the client name you create on Panorama.
7. Click **Approve**.

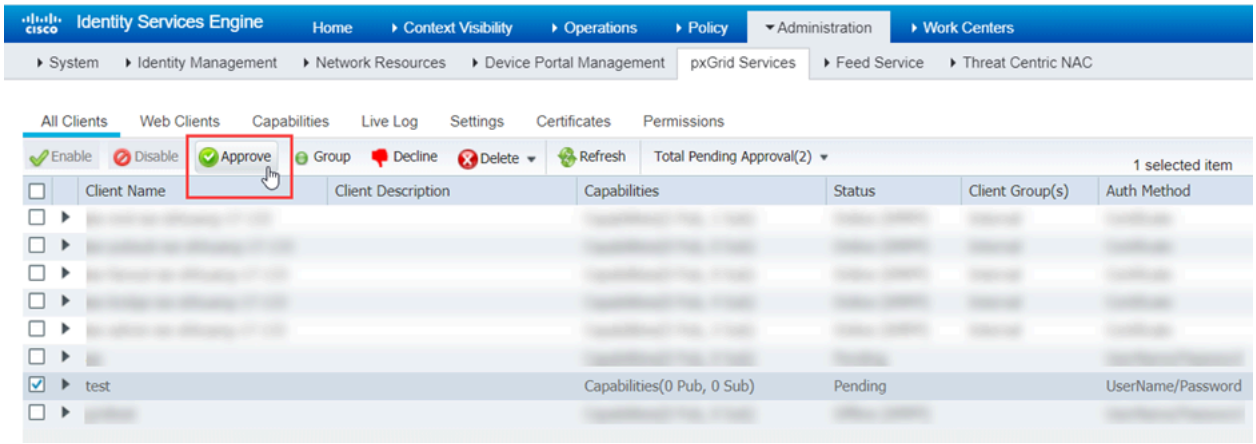

STEP 8 | Add pxGrid server information. The Panorama plugin for Cisco TrustSec supports up to 16 pxGrid (Cisco ISE) servers.

1. Select **Panorama** > **Cisco TrustSec** > **Setup** > **pxGrid Server**.
2. Enter a descriptive **Name** for your pxGrid server.
3. In the **Host** field, enter the IP address or FQDN for your pxGrid server.
4. Enter the client name you created in the previous step.
5. Enter and confirm the client password you generated in the previous step.
6. Verify the pxGrid server identity.
   1. Select **Verify server certificate**.
   2. Select your certificate profile from the **Cert Profile** drop-down.
7. Click **OK**.

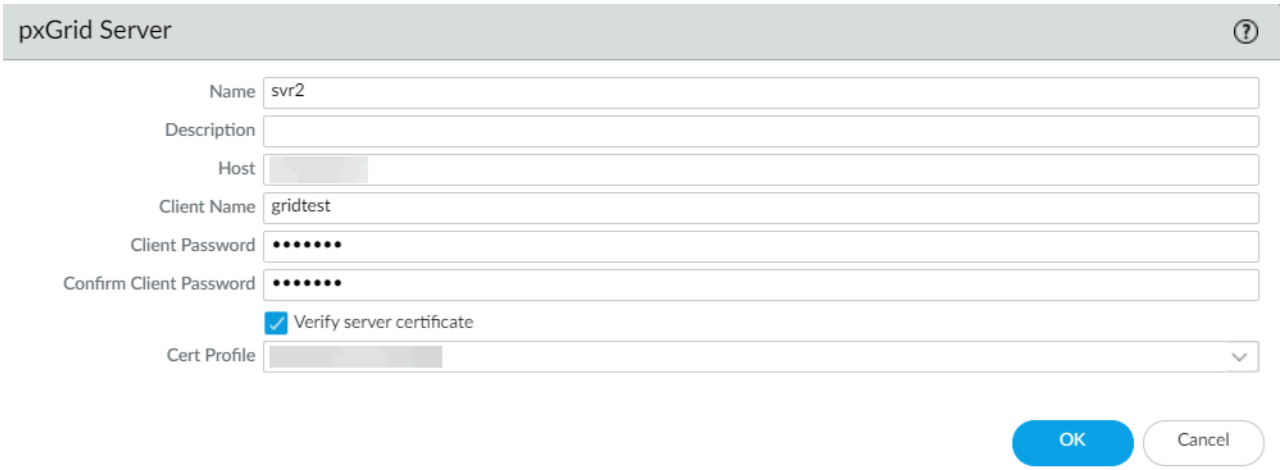

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Configure the Monitoring Definition.

1. Select **Panorama** > **Cisco TrustSec** > **Monitoring Definition** and click **Add**.
2. Enter a descriptive **Name** and optionally a **Description** to identify the monitoring definition.
3. Select the **pxGrid Server**.
4. (Optional) Set Panorama to **Monitor pxGrid sessions in AUTHENTICATED state**. By default, Panorama retrieves IP-Tag mappings from sessions in the "STARTED" state. ISE sessions have the "STARTED" state when there is a corresponding accounting start packet. If no accounting start packet is present for the session, then the session state is "AUTHENTICATED".
5. Select the **Notify Group**.
6. Click **OK**.

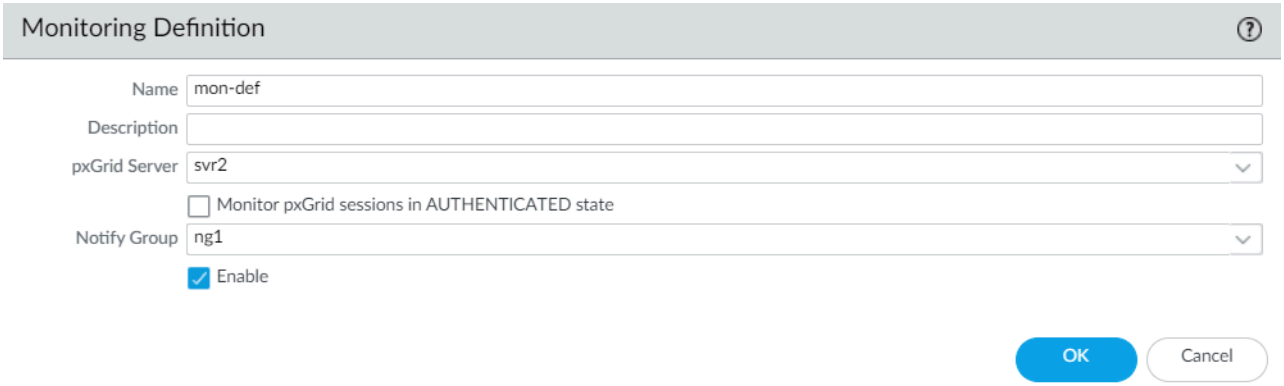

STEP 10 | **Commit** your changes.

STEP 11 | Create active ISE sessions so that Panorama can learn SGT tags for dynamic address group definition. To create active sessions, use ISE to authenticate devices.

Panorama does not collect default SGT tags on ISE.

 ©2024 Palo Alto Networks, Inc.

STEP 12 | Create dynamic address groups and verify that addresses are added to dynamic address groups.

1. Select **Objects** > **Address Groups**.
2. Select the Device Group you created for monitoring endpoints in your Cisco TrustSec environment from the **Device Group** drop-down.
3. Click **Add** and enter a **Name** and **Description** for the dynamic address group.

   The dynamic address group naming convention is: **`cts.svr_<server-name>.sgt_<SGT-name>`**
4. Select **Type** as **Dynamic**.
5. Click **Add Match Criteria**.
6. Select the **And** or **Or** operator and click the plus (+) icon next to the security group name to add it to the dynamic address group.

   Panorama can only display security group tags it has learned from active sessions. Security group tags in live sessions appear in the match criteria list.
7. Select **Panorama** > **Objects** > **Address Groups**.
8. Click **More** in the Addresses column of a dynamic address group.

   Panorama displays a list of IP addresses added to that dynamic address groups based on the match criteria you specified.

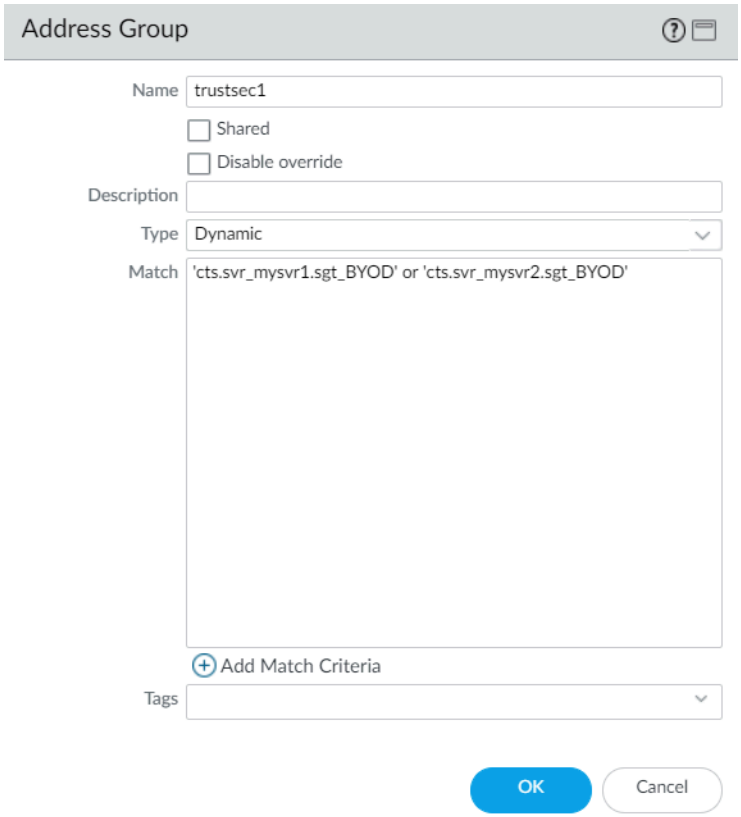

 ©2024 Palo Alto Networks, Inc.

STEP 13 | Use dynamic address groups in policy.

Dynamic address groups are empty until you attach them to a policy. You won't see any IP addresses in your dynamic address group unless a policy is using it.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeats Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

STEP 14 | (Optional) Update the dynamic objects from the pxGrid server at any time by synchronizing dynamic objects. Synchronizing dynamic objects enables you to maintain context on changes in the virtual environment and allows you to enable applications by automatically updating the dynamic address groups used in policy rules.

1. Select **Panorama** > **Cisco TrustSec** > **Monitoring Definition**.
2. Click **Synchronize Dynamic Objects**.

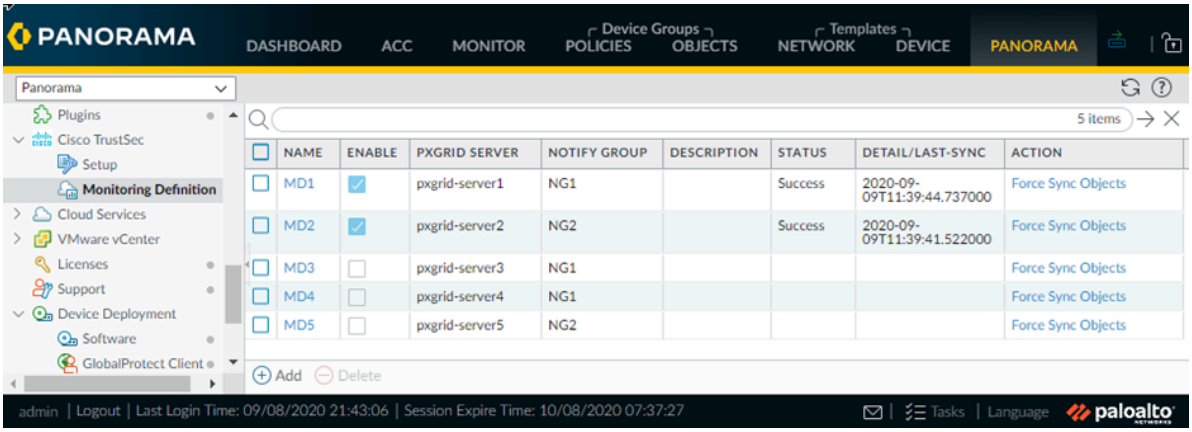

 ©2024 Palo Alto Networks, Inc.

Troubleshoot the Panorama Plugin for Cisco TrustSec

- Plugin Status Commands
- Debug Commands
- Debug Logs

Plugin Status Commands

- Clear counters:

```
clear plugins cisco_trustsec counters
```

- Display monitor status:

```
show plugins cisco_trustsec status
```

- Display counters:

```
show plugins cisco_trustsec counters
```

Debug Commands

- Check IP addresses in dynamic address groups.

```
show object registered-ip tag <tag>
```

```
show object registered-ip all
```

- Fetch the tags of an IP address from a server. The fetched ip-tag mappings are logged in `plugin_cisco_trustsec.log`. No ip-tag mappings are pushed to the notify group associated with the server. No retry if failed.

```
debug plugins cisco_trustsec query pxgrid-server $server-name ip
 $ip-address
```

- Force synchronize with a server and push the mappings to the configd process. No retry if failed.

```
request plugins cisco_trustsec synchronize-dynamic-objects name
 $server-name
```

- Force synchronize with all servers and push the mappings to the configd process. No retry if failed.

```
 request plugins cisco_trustsec synchronize-dynamic-objects all
```

 ©2024 Palo Alto Networks, Inc.

- Force synchronize the mappings from configd process to VM-Series firewalls. No retry if failed.

```
request plugins cisco_trustsec sync
```

Debug Logs

The logs are in the following locations on the disk:

```
/opt/plugins/var/log/pan/plugin_cisco_trustsec.log
/opt/plugins/var/log/pan/plugin_cisco_trustsec_sub.log
/opt/plugins/var/log/pan/plugin_cisco_trustsec_ret.log
/opt/plugins/var/log/pan/plugin_cisco_trustsec_proc.log
```

The size limit for a log file (shared by all plugins installed on your Panorama device) is 10 million bytes. A log file can accept 93,000 session logins. If you configure log rotation, a backup log can support 186,000 session logins.

- Change the plugin debug level.

```
request plugins debug level $level plugin-name cisco_trustsec
```

  - **off**: No debug log.
  - **low**: Dump only basic debug logs.
  - **medium**: Dump detailed debug logs.
  - **high**: Dump everything including request/response messages with servers.
- Merge the logs into a single log file:

```
request plugins cisco_trustsec merge-logs
```

- Show the debug log in the CLI:
  - Cisco TrustSec plugin version 1.0.2 or later installed on a Panorama version earlier than 10.0.0:

```
tail mp-log plugin_cisco_trustsec_merged.log
```

  - Cisco TrustSec plugin version 1.0.2 or later installed on Panorama version 10.0.0 or later:

```
tail follow yes plugins-log
```

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Set Up the VM-Series Firewall on Nutanix AHV

The VM-Series firewall for Nutanix AHV allows you to deploy the VM-Series firewall on devices capable of running the Nutanix Acropolis Hypervisor. If you are using Panorama to manage your VM-Series firewalls on Nutanix AHV, you can use the Panorama plugin for Nutanix to perform VM monitoring. This allows you to dynamically inform the firewall of changes in your Nutanix environment and ensure that policy is applied to virtual machines as they join your network.

- Deploy the VM-Series Firewall on Nutanix AHV
- VM Monitoring on Nutanix

VM Monitoring on Nutanix

Install and configure the Panorama plugin for Nutanix to monitor changes in your Nutanix environment and build policy using dynamic address groups.

- About VM Monitoring on Nutanix
- Install the Panorama Plugin for Nutanix
- Configure the Panorama Plugin for Nutanix

About VM Monitoring on Nutanix

The Panorama plugin for Nutanix facilitates the use of dynamic address groups by monitoring virtual machines in your Nutanix environment. Prism Central groups entities in your Nutanix environments by categories and filters them further by value. Panorama creates tags based on categories and values you define in Prism Central. When a virtual machine is placed in a category and assigned a value, Panorama applies the corresponding tag to the virtual machine's IP address. You can then create security policy by using the tags as match criteria for dynamic address groups in Panorama.

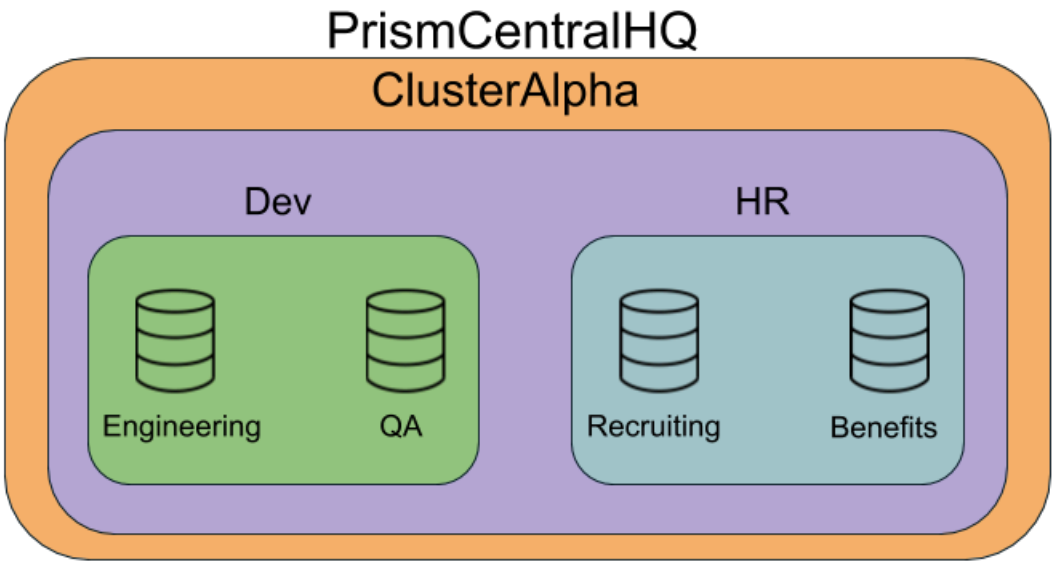

In the example above, we have two categories—Dev and HR—with two values within each of them. And these categories are within the cluster, which is within Prism Central. After you begin monitoring your Nutanix environment, Panorama uses value, category, cluster, and Prism Central to form tags. When you view the match criteria for dynamic address groups, the tags are listed in the following format.

ntnx.PC-<prism-central-name>.CL-<cluster-name>.<category>.<value>

With the information in the example above, Panorama creates the following tags.

`ntnx.PC-PrismCentralHQ.CL-ClusterAlpha.Dev.Engineering`

`ntnx.PC-PrismCentralHQ.CL-ClusterAlpha.Dev.QA`

`ntnx.PC-PrismCentralHQ.CL-ClusterAlpha.HR.Recruiting`

 ©2024 Palo Alto Networks, Inc.

`ntnx.PC-PrismCentralHQ.CL-ClusterAlpha.HR.Benefits`

To secure these workloads in these categories, use tags such as these as match criteria in the dynamic address groups. You can then use the dynamic address groups as source and destination address groups in your security policy rules. When a virtual machine joins a dynamic address group, the policy your created is applied automatically.

Install the Panorama Plugin for Nutanix

To get started with endpoint monitoring on Nutanix, download and install the Panorama plugin for Nutanix.

If you have a Panorama HA configuration, repeat this installation process on each Panorama peer. When installing the plugin on Panoramas in an HA pair, install the plugin on the passive peer before the active peer. After installing the plugin on the passive peer, it will transition to a non-functional state. Installing the plugin on the active peer returns the passive peer to a functional state.

If you have a standalone Panorama or two Panorama appliances installed in an HA pair with multiple plugins installed, plugins might not receive updated IP-tag information if one or more of the plugins is not configured. This occurs because Panorama will not forward IP-tag information to unconfigured plugins. Additionally, this issue can occur if one or more of the Panorama plugins is not in the Registered or Success state (positive state differs on each plugin). Ensure that your plugins are in the positive state before continuing or executing the commands described below.

If you encounter this issue, there are two workarounds:

- Uninstall the unconfigured plugin or plugins. It is recommended that you do not install a plugin that you do not plan to configure right away
- You can use the following commands to work around this issue. Execute the following command for each unconfigured plugin on each Panorama instance to prevent Panorama from waiting to send updates. If you do not, your firewalls may lose some IP-tag information.

  `request plugins dau plugin-name <plugin-name> unblock-device-push yes`

  You can cancel this command by executing:

  `request plugins dau plugin-name <plugin-name> unblock-device-push no`

The commands described are not persistent across reboots and must be used again for any subsequent reboots. For Panorama in HA pair, the commands must be executed on each Panorama.

STEP 1 | Log in to the Panorama user interface.

STEP 2 | Select **Panorama** > **Plugins**.

STEP 3 | Select **Check Now** to retrieve a list of available updates.

STEP 4 | Select **Download** in the Action column to download the plugin.

STEP 5 | Select the version of the plugin and click **Install** in the Action column to install the plugin. Panorama will alert you when the installation is complete.

 ©2024 Palo Alto Networks, Inc.

Configure the Panorama Plugin for Nutanix

After installing the plugin, complete the following procedure to establish a connection between Panorama and Prism Central.

STEP 1 | Log in to the Panorama web interface.

STEP 2 | Enable monitoring and set the monitoring interval.

1. Select **Panorama** > **Nutanix** > **Setup** > **General**.
2. Select **Enable Monitoring**.
3. Set the **Monitoring Interval** in seconds. The monitoring interval is how often Panorama retrieves updated networking information from Prism Central.

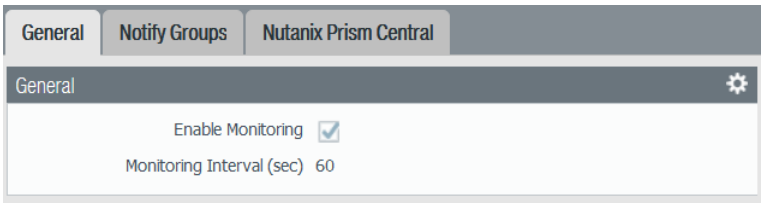

STEP 3 | Create a notify group.

1. Select **Panorama** > **Nutanix** > **Setup** > **Notify Groups**.
2. Click **Add**.
3. Enter a descriptive **Name** for your notify group.
4. Select the device groups in your Nutanix deployment.

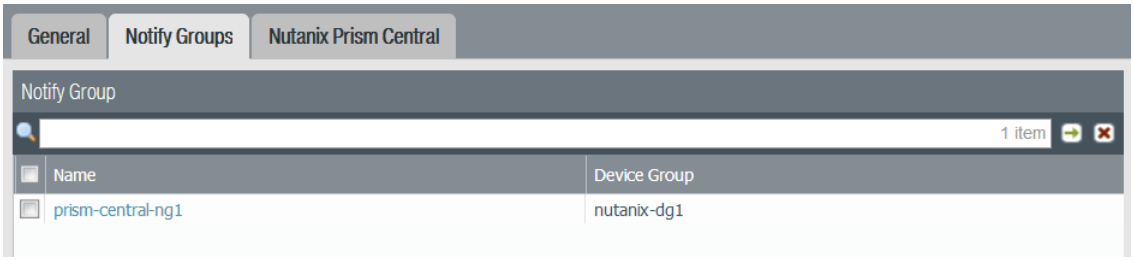

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Add Prism Central information.

1. Select **Panorama** > **Nutanix** > **Setup** > **Nutanix Prism Central**.
2. Click **Add**.
3. Enter a descriptive **Name** for your Prism Central.
4. Enter the IP address or FQDN for Prism Central.
5. Enter your Prism Central username.
6. Enter and confirm your Prism Central password.
7. Click **Validate** to confirm that you entered the Prism Central credentials correctly.

   If you return to the Nutanix Prism Central Info window after clicking OK, clicking the Validate button returns a credential validation error message. This is the expected behavior. Although Panorama displays dots in the password field, the field is empty; this causes the validation to fail despite Panorama being successfully connected to Prism Central.
8. Click **OK**.

STEP 5 | Configure the Monitoring Definition.

1. Select **Panorama** > **Nutanix** > **Monitoring Definition** and click **Add**.
2. Enter a descriptive **Name** and optionally a description to identify the Prism Central for which you use this definition.
3. Select the **Prism Central** and **Notify Group**.
4. Click **OK**.

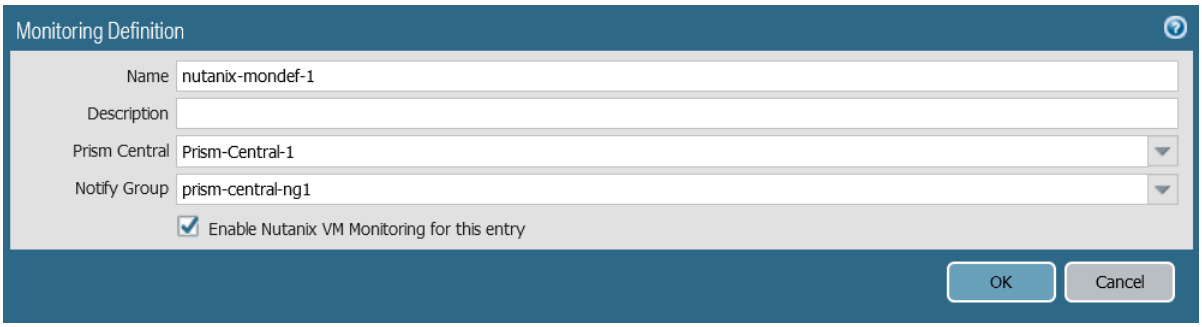

STEP 6 | **Commit** your changes.

 ©2024 Palo Alto Networks, Inc.

STEP 7 | Verify that you can view the VM information on Panorama, and define the match criteria for dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups** and click **Add**.
2. Enter a descriptive **Name** for your dynamic address groups.
3. Select **Dynamic** from the Type drop-down.
4. Click **Add Match Criteria**. You can select dynamic tags as the match criteria to populate the members of the group. Select the **And** or **Or** operator and select the attributes that you would like to filter for or match against. and then click **OK**.
5. **Commit** your changes.

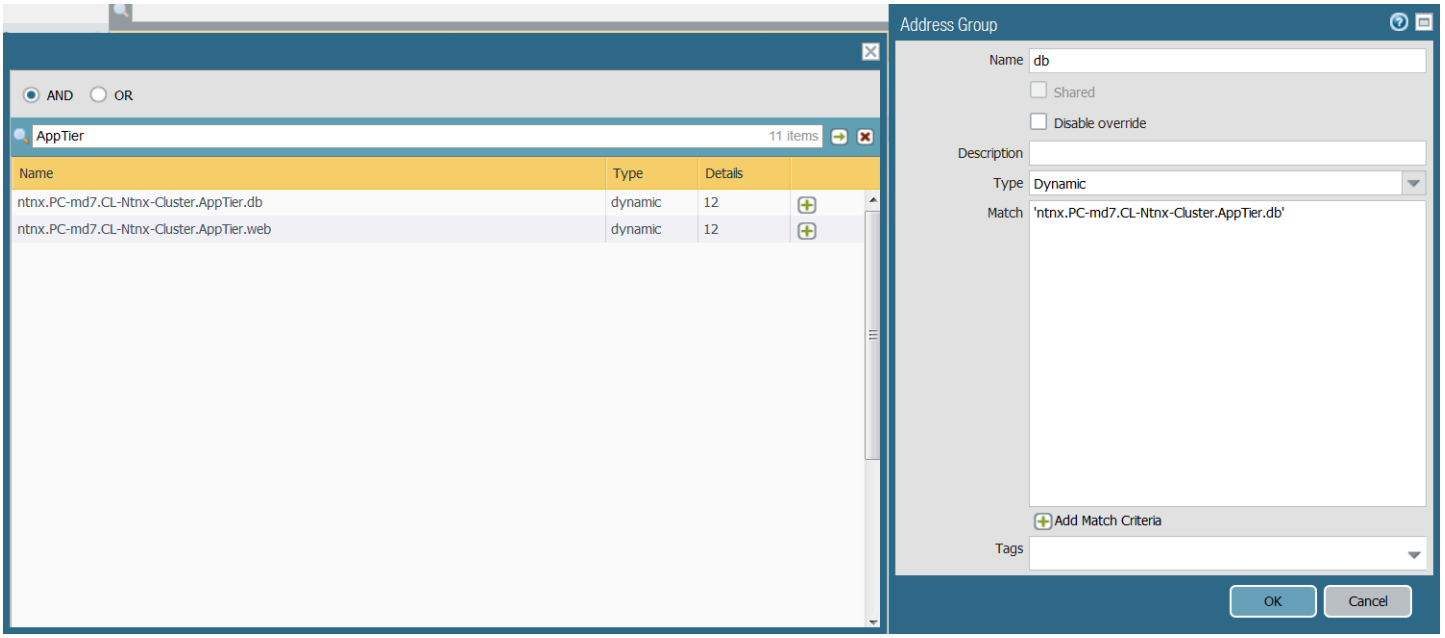

STEP 8 | Verify that addresses in your VMs are added to dynamic address groups.

1. Select **Panorama** > **Objects** > **Address Groups**.
2. Click **More** in the Addresses column of a dynamic address group.

   Panorama displays a list of IP addresses added to that dynamic address group based on the match criteria you specified.

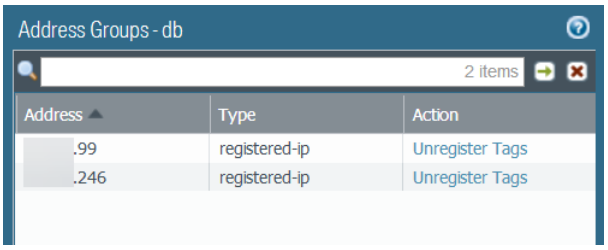

 ©2024 Palo Alto Networks, Inc.

STEP 9 | Use dynamic address groups in policy.

1. Select **Policies** > **Security**.
2. Click **Add** and enter a **Name** and a **Description** for the policy.
3. Add the **Source Zone** to specify the zone from which the traffic originates.
4. Add the **Destination Zone** at which the traffic is terminating.
5. For the **Destination Address**, select the Dynamic address group you just created.
6. Specify the action— **Allow** or **Deny**—for the traffic, and optionally attach the default security profiles to the rule.
7. Repeat Steps 1 through 6 to create another policy rule.
8. Click **Commit**.

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall

Bootstrapping allows you to create a repeatable and streamlined process of deploying new VM-Series firewalls on your network because it allows you to create a package with the model configuration for your network and then use that package to deploy VM-Series firewalls anywhere.

You can either bootstrap the firewall with **complete** configuration so that the firewall is fully configured at startup, or you can begin with a **basic** configuration—a minimal initial configuration that enables you to boot the firewall and then register with Panorama to complete the configuration.

If you choose the **basic** configuration and you are deploying on AWS, Azure or GCP, you can use the bootstrap package and an `init-cfg.txt` file. Alternatively, you can bootstrap with **user data**. Instead of providing bootstrap configuration parameters in files, you enter them as key-value pairs directly into the AWS or GCP user interface when you launch a VM-Series firewall. Azure has a similar process with which you provide the bootstrap parameters in a template or other text file accessed from the Azure CLI.

If you create the bootstrap package, you deliver it from an external device (such as a virtual disk, a virtual CD-ROM, or a cloud storage device (such as a bucket).

- Choose a Bootstrap Method
- VM-Series Firewall Bootstrap Workflow
- Bootstrap Package
- Bootstrap Configuration Files
- Generate the VM Auth Key on Panorama
- Create the init-cfg.txt File
- Create the bootstrap.xml File
- Prepare the Licenses for Bootstrapping
- Prepare the Bootstrap Package
- Bootstrap the VM-Series Firewall on AWS
- Bootstrap the VM-Series Firewall on Azure
- Bootstrap the VM-Series Firewall on ESXi
- Bootstrap the VM-Series Firewall on Google Cloud Platform
- Bootstrap the VM-Series Firewall on Hyper-V
- Bootstrap the VM-Series Firewall on KVM
- Verify Bootstrap Completion
- Bootstrap Errors

Choose a Bootstrap Method

You can bootstrap the VM-Series firewall with a **basic configuration** or a **complete configuration**.

A **complete** configuration uses the bootstrap package and includes everything you need to fully configure the firewall at boot up. This includes configuration parameters (in init-cfg.txt), content updates, and software versions. A complete configuration can include both init-cfg.txt and bootstrap.xml files.

Configuration Method	Configuration Location	Comment
Specify complete configuration information in /config/bootstrap.xml in the bootstrap package.	**Public cloud storage** AWS S3 bucket, Azure storage account, or Google storage bucket.	• Full bootstrap package in the storage bucket. • Requires cloud storage and an IAM role to access it.

A **basic** configuration is a minimal configuration that enables you to launch, license, and register the VM-Series firewall. The basic configuration does not support plugins, content, software images, or bootstrap.xml.

After you boot the firewall you can connect with Panorama to complete the configuration, or log in to the firewall to update content and software manually. The following table briefly contrasts three ways you can store and access a basic configuration:

Configuration Method	Configuration Location	Comment
init-cfg.txt Store basic configuration parameters as key-value pairs in config/init-cfg.txt in the bootstrap package.	**Public cloud storage** • AWS S3 bucket • Azure storage account • GCP storage bucket	• Requires cloud storage and an IAM role to access it. The Panorama admin must also be granted access to the bucket.
User data Enter configuration parameters into the public cloud user interface as key-value pairs.	**VM Instance** • Alibaba: User data • AWS: User data • Azure: Custom data • GCP: GCP metadata • Oracle Cloud Infrastructure: User data	• The initial configuration parameters are stored with the VM. • No need for separate storage and the associated IAM role.
AWS Secret Manager Enter configuration parameters into the AWS secret manager as key-value pairs.	Encrypted in AWS Secret Manager.	• You need an IAM role to create a secret. Others can be granted permission to get the secret.

 ©2024 Palo Alto Networks, Inc.

Configuration Method	Configuration Location	Comment
		• To get the secret, pass the secret name using user data.

See the VM-Series firewall bootstrap workflow to compare the workflow for the basic and complete configurations.

- Basic Configuration
- Complete Configuration

Basic Configuration

A basic configuration includes the initial configuration and licenses. You can use the bootstrap package to pass the key-value pairs for the initial configuration, or you can enter the bootstrap parameters key-value pairs as user data.

If you do not use Panorama, you can use the initial configuration to bootstrap the firewall, then log in and complete the configuration manually. If you use Panorama, your initial configuration must include bootstrap parameters for the IP addresses for your Panorama servers and the VM Auth Key so the bootstrapped firewall can register with Panorama and complete the full configuration.

- Add a Basic Configuration to the Bootstrap Package
- Enter a Basic Configuration as User Data (Public Clouds)
- Save a Basic Configuration in the AWS Secrets Manager

Add a Basic Configuration to the Bootstrap Package

The initial configuration is a minimal configuration that enables you to launch, license, and register the VM-Series firewall, and connect with Panorama, if applicable. You deliver the configuration (init-cfg.txt) in the bootstrap package.

Enter a Basic Configuration as User Data (Public Clouds)

When you deploy the VM-Series firewall from a public cloud user interface, you can enter the configuration parameters as user data during the launch/deployment process. If you have sufficient permissions to deploy a firewall from your cloud account, and access Panorama (if you are using it), you can skip creating a bootstrap package, creating configuration files, and other bootstrap tasks related to cloud storage (a storage bucket, IAM roles, or service accounts that grant external access to storage).

Configuration parameters include the values in init-cfg.txt File Components, and the following additional values only available as user data:

- **`authcodes`**—The authcode use to register the VM-Series firewall. For example, **`authcodes=I7115398`**.
- **`mgmt-interface-swap`**—Used to swap the management interface when the VM-Series firewall is behind a load balancer in an AWS or GCP deployment. For example, **`mgmt-interface-swap=enable`**.

 ©2024 Palo Alto Networks, Inc.

You can enter configuration parameters as key-value pairs directly into the Alibaba, AWS, GCP or OCI user interface. You can also define the configuration from a text file or a cloud-native template, such as an AWS Cloud Formation template, Azure ARM template, a GCP YAML file, or a Terraform template.

Each cloud has a different term for user data, and uses different separators between bootstrap parameters.

- **Alibaba Cloud User Data** —Use a newline (\n) for each parameter, and if a parameter has multiple options, use commas to separate them.
- **AWS User Data**—Use a semicolon or newline (\n). If a parameter has more than one option, separate options with a comma. For example:

  If you choose to save your basic configuration in the AWS Secrets Manager, enter the secret name as a key-value pair in the user data field. For example:

```
type=dhcp-client
hostname=palo1
panorama-server=<PANORAMA-1 IP>
panorama-server-2=<PANORAMA-2 IP>
tplname=STK-NGFW-01
dgname=DG-NGFW-01
dns-primary=169.254.169.253
dns-secondary=8.8.8.8
op-command-modes=mgmt-interface-swap
dhcp-send-hostname=yes
dhcp-send-client-id=yes
dhcp-accept-server-hostname=yes
dhcp-accept-server-domain=yes
vm-auth-key= <YOUR AUTH KEY HERE>
authcodes= <<YOUR AUTH CODE HERE>
```

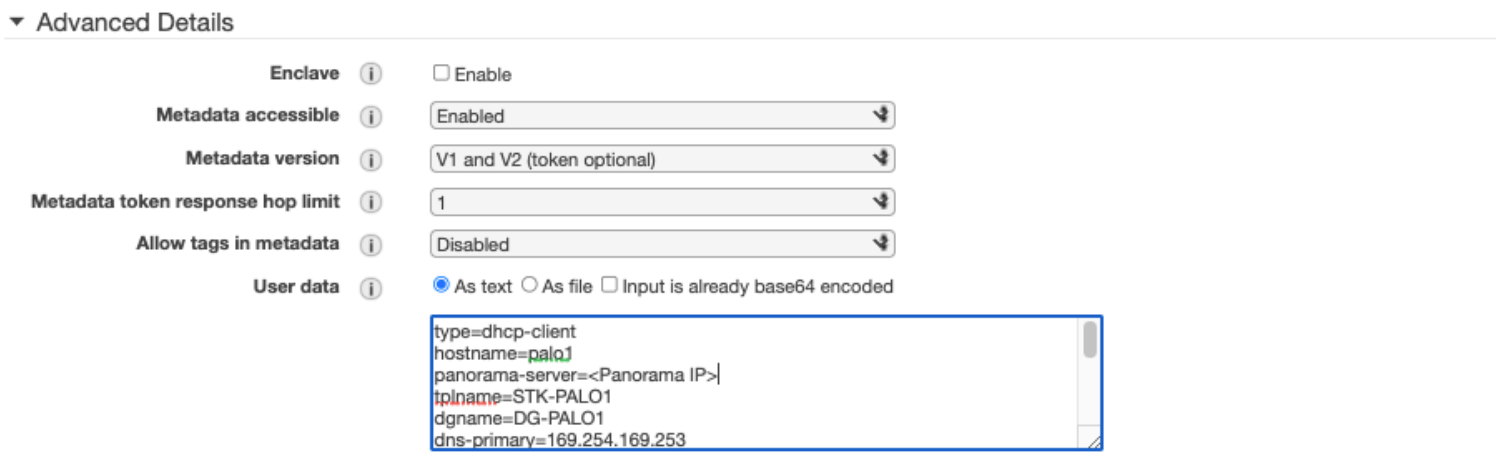

 ©2024 Palo Alto Networks, Inc.

- **Azure Custom Data**—Use a semicolon. If a parameter has more than one option, separate options with a comma. For example:

```
type=dhcp-client; op-command-modes=jumbo-frame;
vm-series-auto-registration-pin-id=abcdefgh1234****;
vm-series-auto-registration-pin-value=zyxwvut-0987****
```

- **GCP Custom Metadata**—In a file, such as a YAML file or Terraform template, use a newline (\n) for each parameter, and if a parameter has multiple options, use commas to separate them. For example:

```
type=dhcp-client
op-command-modes=mgmt-interface-swap,jumbo-frame
vm-series-auto-registration-pin-id=abcdefgh1234****
vm-series-auto-registration-pin-value=zyxwvut-0987****
```

- **Oracle Cloud Infrastructure User Data**—Use a newline (\n) for each parameter, and if a parameter has multiple options, use commas to separate them.

Save a Basic Configuration in the AWS Secrets Manager

You can use the AWS Secrets Manager to store the basic configuration as a secret, and then use User Data to bootstrap a VM with the parameters stored in the secret. To perform this task you need permission to use the Secrets Manager.

- The secret creator must have full Secrets Manager administrator permissions. A Secrets Manager administrator can permit others to use the secret, as described in Authentication and access control for AWS Secrets Manager.

  For example, the following policy statement allows you to get the secret value:

```
{
  "Version": "2012-10-17",
    {
      "Effect": "Allow",
      "Action": "secretsmanager:GetSecretValue",
      "Resource":
      "arn:aws:secretsmanager:us-east-1:688382******:secret:My_bts-
******"
    }
}
```

  Refer to Actions, Resources, and Context Keys You Can Use in an IAM Policy or Secret Policy for AWS Secrets Manager to see actions that require permission, such as list, get, and rotate secret.

- (Optional) To encrypt the secret you can use the DefaultEncryptionKey from AWS Secrets Manager.

STEP 1 | Log in to the AWS console and under Security, Identity and Compliance, select **Secrets Manager** and select **Store a new secret**.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Select **other type of secrets**.

1. Enter the key-value pairs to define the basic configuration.

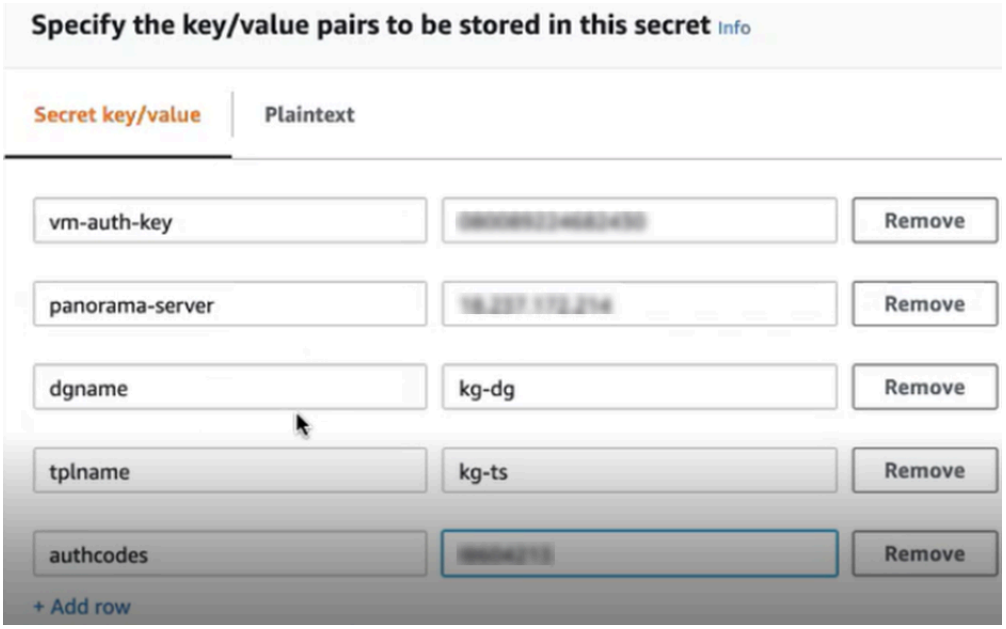

> ***mgmt-interface-swap*** *does not work as a key-value pair in an AWS secret. It must be entered as:* ***op-command-modes=mgmt-interface-swap***

2. Select the DefaultEncryptionKey, and click **Next**.

STEP 3 | Supply the secret name and description.

1. Edit the resource permissions to securely access secrets across AWS accounts. For example:

```
{
    "Version": "2012-10-17",
    "Statement": [
    {
        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": "s3:ListBucket",
        "Resource": "arn:aws:s3:::sn-bootstrap"
    },
    {
        "Sid": "VisualEditor1",
        "Effect": "Allow",
        "Action": "s3:GetObject",
        "Resource": "arn:aws:s3:::sn-bootstrap/*"
    },
    {
        "Effect": "Allow",
        "Action": "secretsmanager:GetSecretValue",
        "Resource": "arn:aws:secretsmanager:us-
east-1:688382******:
                            secret:My_bootstrap"
    }
```

 ©2024 Palo Alto Networks, Inc.

```
  ]
}
```

2. (Optional) You can examine the secret from the command line (if you have permission). For example:

```
# aws secretsmanager get-secret-value --secret-id My_bootstrap
{
    "ARN": "arn:aws:secretsmanager:us-east-1:688382******:
                secret:My_bootstrap",
    "Name": "My_bootstrap",
    "VersionId": "01b6853d-e187-479f-**********",
    "SecretString": "{\"mgmt-interface-swap\":\"enable\",
        \"vm-auth-key\":\"AAA\",\"panorama-server\":
\"10.*.*.1\",
        \"panorama-server-2\":\"10.*.*.2\",\"dgname\":\"dg-
s0000h\",
        \"tplname\":\"tpl-santosh\",\"license-authcode\":
\"AAAA\"}",
    "VersionStages": [ "AWSCURRENT" ],
    "CreatedDate": 1581018411.847
}
```

Complete Configuration

A complete configuration ensures the firewall is fully configured on boot up. The bootstrap.xml file includes the initial configuration, licenses, software, content, and a version of the VM-Series plugin. You can create `bootstrap.xml` manually or you can export an existing configuration, as described in Create the bootstrap.xml File.

 ©2024 Palo Alto Networks, Inc.

VM-Series Firewall Bootstrap Workflow

Use the following workflow to bootstrap your VM-Series firewall. Refer to the following figure for an overview of both complete and basic bootstrapping procedures.

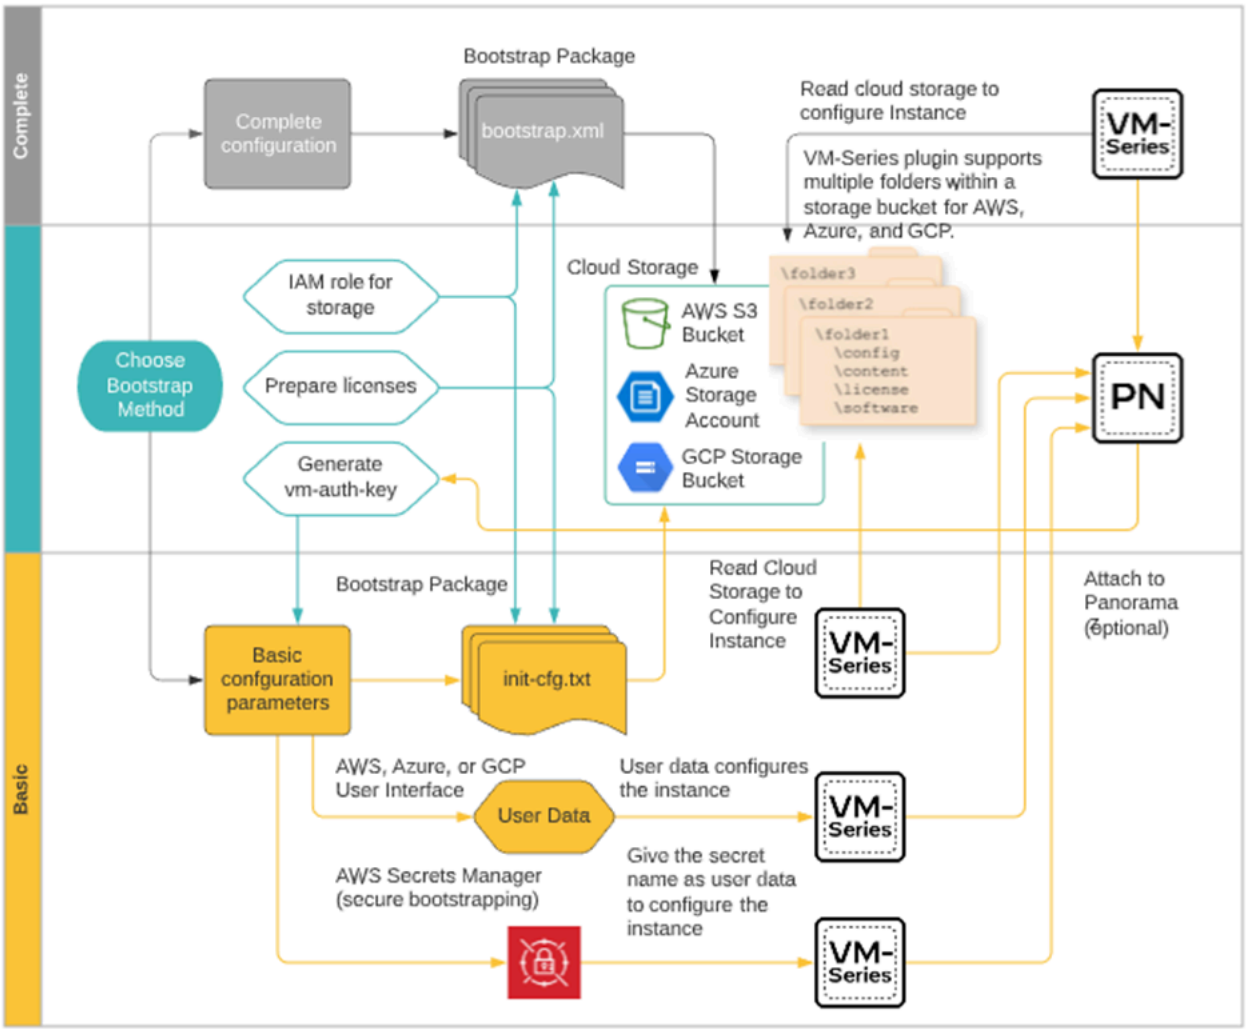

STEP 1 | (Optional) For security reasons, you can only bootstrap a firewall when it is in factory default state. If you want to use the bootstrap package to bootstrap a VM-Series firewall that has been previously configured, reset the firewall to factory default settings.

STEP 2 | Choose a bootstrap method.

After you familiarize yourself with the bootstrap package, assess whether you want to use a complete configuration, or use a basic configuration and optionally use Panorama to manage the bootstrapped firewall.

If you choose a basic configuration, decide whether to use the bootstrap package, or enter the configuration parameters as key-value pairs in user data.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | (Optional) If you want to use Panorama to manage the VM-Series firewalls being bootstrapped, generate the VM auth key on Panorama. You must include this key in the `init-cfg.txt` file (**`vm-auth-key`**) or enter the key-value pair as user data.

STEP 4 | Prepare the licenses for bootstrapping.

STEP 5 | If you choose the basic configuration and you plan to bootstrap with user data, skip to Step 7.

If you plan to use the basic configuration and the bootstrap package, create the init-cfg.txt file and prepare the bootstrap package.

If you choose the complete configuration, create the bootstrap.xml file and prepare the full bootstrap package.

STEP 6 | Prepare the bootstrap package and save the bootstrap package in the appropriate delivery format for your hypervisor.

STEP 7 | Bootstrap the VM-Series firewall.

- Bootstrap the VM-Series Firewall on AWS
- Bootstrap the VM-Series Firewall on Azure
- Bootstrap the VM-Series Firewall on ESXi
- Bootstrap the VM-Series Firewall on Google Cloud Platform
- Bootstrap the VM-Series Firewall on Hyper-V
- Bootstrap the VM-Series Firewall on KVM

STEP 8 | Verify bootstrap completion.

 ©2024 Palo Alto Networks, Inc.

Bootstrap Package

The bootstrap process is initiated only on first boot when the firewall is in a factory default state.

- Bootstrap Package Structure
- Bootstrap Package Delivery

Bootstrap Package Structure

The bootstrap package must include the `/config`, `/license`, `/software`, and `/content` folders, even if they are empty. The `/plugins` folder is optional. For an example, see Prepare the Bootstrap Package.

- **/config folder**—Contains the configuration files. The folder can hold two files: `init-cfg.txt` and the `bootstrap.xml`. For details, see Bootstrap Configuration Files.

  If you intend to pre-register VM-Series firewalls with Panorama with bootstrapping, you must generate a VM auth key on Panorama and include the generated key in the `init-cfg.txt` file. See Generate the VM Auth Key on Panorama.

- **/license folder**—Contains the license keys or auth codes for the licenses and subscriptions that you intend to activate on the firewalls. If the firewall does not have Internet connectivity, you must either manually obtain the license keys from the Palo Alto Networks Support portal or use the Licensing API to obtain the keys and then save each key in this folder. For details, see Prepare the Licenses for Bootstrapping.

  You must include an auth code bundle instead of individual auth codes so that the firewall or orchestration service can simultaneously fetch all license keys associated with a firewall. If you use individual auth codes instead of a bundle, the firewall will retrieve only the license key for the first auth code included in the file.

- **/software folder**—Contains the software images required to upgrade a newly provisioned VM-Series firewall to the desired PAN-OS version for your network. You must include all intermediate software versions between the current version and the final PAN-OS software version to which you want to upgrade the VM-Series firewall. Refer to VM-Series Firewall Hypervisor Support in the Compatibility Matrix.
- **/content folder**—Contains the application and threat updates, WildFire updates, and the BrightCloud URL filtering database for the valid subscriptions on the VM-Series firewall. You must include the minimum content versions required for the desired PAN-OS version. If you do not have the minimum required content version associated with the PAN-OS version, the VM-Series firewall cannot complete the software upgrade.
- **/plugins folder**—Optional folder contains a single VM-Series plugin image.

Bootstrap Package Delivery

The file type used to deliver the bootstrap package to the VM-Series firewall varies based on your hypervisor. Use the table below to determine the file type your hypervisor or cloud vendor supports.

 ©2024 Palo Alto Networks, Inc.

External Device for Bootstrapping (Bootstrap Package Format)	AWS	Azure	ESXi	Google	Hyper-V	KVM
CD-ROM (ISO image)	—	—	Yes	—	Yes	Yes
Block Storage Device	—	—	Yes	—	Yes	Yes
Storage Account	—	Yes	—	—	—	—
Storage Bucket	Yes	—	—	Yes	—	—

When you attach the storage device to the firewall, the firewall scans for a bootstrap package and, if one exists, the firewall uses the settings defined in the bootstrap package.

If you have included a Panorama server IP address in the file, the firewall connects with Panorama. If the firewall has Internet connectivity, it contacts the licensing server to update the UUID and obtain the license keys and subscriptions. The firewall is then added as an asset in the Palo Alto Networks Support portal. If the firewall does not have Internet connectivity, it either uses the license keys you included in the bootstrap package, or it connects to Panorama to retrieve the appropriate licenses and deploys them to the managed firewalls.

 ©2024 Palo Alto Networks, Inc.

Bootstrap Configuration Files

The bootstrap package must include the basic configuration in `config/init-cfg.txt`. The complete configuration (in `/config/bootstrap.xml` file) is optional.

When you include `init-cfg.txt` file and the `bootstrap.xml` file in the bootstrap package, the firewall merges the configurations of those files, and if any settings overlap, the firewall uses the values defined in the `init-cfg.txt` file.

- init-cfg.txt
- bootstrap.xml

init-cfg.txt

Contains basic information for configuring the management interface on the firewall, such as the IP address type (static or DHCP), IP address (IPv4 only or both IPv4 and IPv6), netmask, and default gateway. The DNS server IP address, Panorama IP address and device group and template stack parameters are optional.

You can use the generic name `init-cfg.txt`, or to be more specific, you can prepend the UUID or Serial number of each firewall to the filename (for example: `0008C100105-init-cfg.txt`).

When the firewall boots, it searches for a text file that matches its UUID or serial number and, if none is found, it searches using the generic filename `init-cfg.txt`. For a sample file, see Create the init-cfg.txt File.

If you are using Panorama to manage your bootstrapped VM-Series firewalls:

- *You must generate a VM auth key on Panorama and include the key in the `init-cfg.txt` file. For more information, see Generate the VM Auth Key on Panorama.*
- *The Panorama appliance that manages the firewalls must be in Panorama mode. If you use a Panorama appliance in Management-Only mode, firewall logs are dropped because Panorama in Management-Only mode does not have a Log Collector Group that can store firewall logs.*

bootstrap.xml

The optional `bootstrap.xml` file contains a complete configuration for the firewall. If you are not using Panorama to centrally manage your firewalls, the `bootstrap.xml` file provides a way to automate the process of deploying firewalls that are configured at launch.

You can define the configuration manually or export the running configuration (`running-config.xml`) from an existing firewall and save the file as `bootstrap.xml`. If you export `bootstrap.xml` file, make sure to export the XML file from a firewall deployed on the same platform or hypervisor as your deployment. See Create the bootstrap.xml File.

 ©2024 Palo Alto Networks, Inc.

Generate the VM Auth Key on Panorama

If you want to use Panorama to manage the VM-Series firewalls that you are bootstrapping, you must generate a VM auth key on Panorama and include the key in the basic configuration (init-cfg.txt) file. The VM auth key allows Panorama to authenticate the newly bootstrapped VM-Series firewall. So, to manage the firewall using Panorama, you must include the IP address for Panorama and the VM auth key in the basic configuration file as well as the license auth codes in the /license folder of the bootstrap package. The firewall can then provide the IP address, serial number, and the VM auth key in its initial connection request to Panorama so that Panorama can verify the validity of the VM auth key and add the firewall as a managed device. If you provide a device group and template in the basic configuration file, Panorama will assign the firewall to the appropriate device group and template so that you can centrally configure and administer the firewall using Panorama.

The lifetime of the key can vary between 1 hour and 8760 hours (1 year). After the specified time, the key expires and Panorama will not register VM-Series firewalls without a valid auth-key in this connection request.

STEP 1 | Log in to the Panorama CLI or access the API:

- In the CLI, use the following operational command:

  ```
  request bootstrap vm-auth-key generate lifetime <1-8760>
  ```

  For example to generate a key that is valid for 24 hrs, enter the following:

  ```
  request bootstrap vm-auth-key generate lifetime 24
  VM auth key 755036225328715 generated. Expires at: 2015/12/29
   12:03:52
  ```

- In the API, use the following URL:

  ```
  https://<Panorama_IP_address>/api/?
  type=op&cmd=<request><bootstrap><vm-auth-
  key><generate><lifetime><number-of-hours></lifetime></generate></
  vm-auth-key></bootstrap></request>
  ```

  where the lifetime is the number of hours for which the VM auth key is valid.

 ©2024 Palo Alto Networks, Inc.

STEP 2 | Verify the validity term of the VM auth key(s) you generated on Panorama. Make sure that the validity term allows enough time for the firewall(s) to register with Panorama.

```
https://<Panorama_IP_address>/api/?
type=op&cmd=<request><bootstrap><vm-auth-key><show></show></vm-
auth-key></bootstrap></request>
```

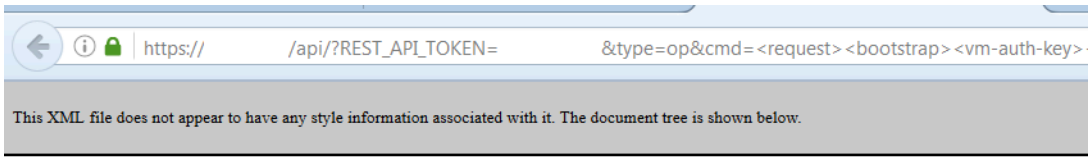

This XML file does not appear to have any style information associated with it. The document tree is shown below.


```
- <response status="success">
  - <result>
    - <bootstrap-vm-auth-keys>
      - <entry>
          <vm-auth-key>085812955845977</vm-auth-key>
          <expiry-time>2016/03/17 08:35:05</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>136387033275034</vm-auth-key>
          <expiry-time>2016/05/20 14:12:59</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>178644792323541</vm-auth-key>
          <expiry-time>2016/06/10 16:25:36</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>221348425464173</vm-auth-key>
          <expiry-time>2016/05/20 13:54:25</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>245832696687351</vm-auth-key>
          <expiry-time>2015/12/22 17:53:48</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>386239691539160</vm-auth-key>
          <expiry-time>2016/03/02 11:09:46</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>420246530153909</vm-auth-key>
          <expiry-time>2016/03/09 00:57:01</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>431216710324086</vm-auth-key>
          <expiry-time>2016/03/09 00:57:09</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>445486056501180</vm-auth-key>
          <expiry-time>2016/05/20 14:12:52</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>633795692572911</vm-auth-key>
          <expiry-time>2016/03/09 14:50:38</expiry-time>
        </entry>
      - <entry>
          <vm-auth-key>798346857952985</vm-auth-key>
          <expiry-time>2016/05/20 14:08:14</expiry-time>
        </entry>
      </bootstrap-vm-auth-keys>
    </result>
  </response>
```


STEP 3 | Add the generated VM auth key to the basic configuration (init-cfg.txt) file. See Create the init-cfg.txt File

 ©2024 Palo Alto Networks, Inc.

Create the init-cfg.txt File

The `init-cfg.txt` file is required to bootstrap the VM-Series firewall. It provides the basic information the firewall needs to connect to your network.

- init-cfg.txt File Components
- Sample init-cfg.txt File

Complete the following procedure to create the `init-cfg.txt` file.

STEP 1 | Create a new text file.

Use a text editor such as Notepad, EditPad, or other plain-text editors to create a text file.

STEP 2 | Add the basic network configuration for the management interface on the firewall.

If any of the required parameters are missing in the file, the firewall exits the bootstrap process and boots up using the default IP address, 192.168.1.1. You can view the system log on the firewall to detect the reason for the bootstrap failure. For errors, see Licensing API.

There are no spaces between the key and value in each field. Do not add spaces as they could cause failures during parsing on the mgmtsrvr side.

- To configure the management interface with a static IP address, you must specify the IP address, type of address, default gateway, and netmask. An IPv4 address is required, IPv6 address is optional. For syntax, see Sample init-cfg.txt File.
- To configure the management interface as a DHCP client, you must specify only the type of address. If you enable the DHCP client on the management interface, the firewall ignores the IP address, default gateway, netmask, IPv6 address, and IPv6 default gateway values defined in the file. For syntax, see Sample init-cfg.txt File.

When you enable DHCP on the management interface, the firewall takes the DHCP assigned IP address and is accessible over the network. You can view the DHCP assigned IP address on the General Information widget on the Dashboard or with the CLI command **`show system info`**. However, the default static management IP address 192.168.1.1 is retained in the running configuration (**`show config running`**) on the firewall. This static IP address ensures that you can always restore connectivity to your firewall, in the event you lose DHCP access to the firewall.

STEP 3 | Add the VM auth key to register a VM-Series firewall with Panorama.

To add a VM-Series firewall on Panorama, you must add the VM auth key that you generated on Panorama to the basic configuration (`init-cfg.txt`) file. For details on generating a key, see Generate the VM Auth Key on Panorama.

STEP 4 | Add details for accessing Panorama.

- Add IP addresses for the primary and secondary Panorama servers.
- Specify the template and the device group to which you want to assign the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 5 | (Recommended) Add the VM-Series registration pin and value for installing the device certificate.

If you want to install the device certificate on the VM-Series firewall at launch, you must generate the VM-Series registration pin ID and value on the CSP and include it in the `init-cfg.txt` file. This pin and value also applies any site licenses that use the PAYG license.

STEP 6 | (Optional) Include additional parameters for the firewall.

- Add IP address for the primary and secondary DNS servers.
- Add the hostname for the firewall.
- Enable either jumbo frames or multiple-virtual systems (or both)
- Enable management interface swap (mgmt) and the dataplane interface (ethernet 1/1) for the VM-Series firewall on AWS or GCP. For more information on changing the management interface, see Management Interface Mapping for Use with Amazon ELB or Management Interface Swap for Google Cloud Platform Load Balancing.
- Enable or disable DPDK.

init-cfg.txt File Components

The following table describes the bootstrap parameters in the `init-cfg.txt` file.

Field	Description
type=	Type of management IP address: static or dhcp-client. This field is required.
ip-address=	IPv4 address. This field is ignored if the type is dhcp-client. If the type is static, an IPv4 address is required; the ipv6-address field is optional and can be included. You cannot specify the management IP address and netmask configuration for the VM-Series firewall in AWS and Azure. If defined, the firewall ignores the values you specify.
default-gateway=	IPv4 default gateway for the management interface. This field is ignored if the type is dhcp-client. If the type is static, and ip-address is used, this field is required.
netmask=	IPv4 netmask. This field is ignored if the type is dhcp-client. If the type is static, and ip-address is used, this field is required.
ipv6-address=	(Optional) IPv6 address and /prefix length of the management interface. This field is ignored if the type is dhcp-client. If the type is static, this field can be specified along with the ip-address field, which is required.

 ©2024 Palo Alto Networks, Inc.

Field	Description
ipv6-default-gateway=	IPv6 default gateway for the management interface. This field is ignored if the type is dhcp-client. If the type is static and ipv6-address is used, this field is required.
hostname=	Host name for the firewall.
panorama-server=	IPv4 or IPv6 address of the primary Panorama server. This field is not required but recommended for centrally managing your firewalls.
panorama-server-2=	IPv4 or IPv6 address of the secondary Panorama server. This field is not required but recommended.
tplname=	Panorama template stack name. If you add a Panorama server IP address, as a best practice assign the firewall to a template stack on Panorama and enter the template stack name in this field so that you can centrally manage and push configuration settings to the firewall.
dgname=	Panorama device group name. If you add a Panorama server IP address, as a best practice create a device group on Panorama and enter the device group name in this field so that you can group the firewalls logically and push policy rules to the firewall.
cgname=	Panorama collector group name. If you want to bootstrap the firewall to send logs to a Panorama collector group, you must first configure a collector group on Panorama and then configure the firewall to forward logs to Panorama. On the M-Series appliances, a default Collector Group is predefined and already contains the local Log Collector as a member. On the Panorama virtual appliance, you must add the Collector Group and add the local Log Collector as a member.
dns-primary=	IPv4 or IPv6 address of the primary DNS server.
dns-secondary=	IPv4 or IPv6 address of the secondary DNS server.
vm-auth-key=	Virtual machine authentication key for Panorama (see Generate the VM Auth Key on Panorama). This field is ignored when bootstrapping hardware firewalls.
op-command-modes=	The following values are allowed: multi-vsys, jumbo-frame, mgmt-interface-swap. If you enter multiple values, use a space or a comma to separate the entries. • **multi-vsys**—(Hardware-based firewalls only) Enables multiple virtual systems.

 ©2024 Palo Alto Networks, Inc.

Field	Description

	• **jumbo-frame**—Enables the default MTU size for all Layer 3 interfaces to be set at 9192 bytes.
	• **mgmt-interface-swap**—(VM-Series firewall on AWS and Google only) Allows you to swap the management interface (MGT) with the dataplane interface (ethernet 1/1) when deploying the firewall. For details, see:
	  • Management Interface Mapping for Use with Amazon ELB
	  • Management Interface Swap for Google Cloud Platform Load Balancing
	  • Use the VM-Series CLI to Swap the Management Interface on ESXi
	  • Use the VM-Series CLI to Swap the Management Interface on KVM

op-cmd-dpdk-pkt-io=	The value **on** or **off** allows you to enable or disable Data Plane Development Kit (DPDK) in environments where the firewall supports DPDK. DPDK allows the host to process packets faster by bypassing the Linux kernel; interactions with the NIC are performed using drivers and the DPDK libraries.
plugin-op-commands=	Specify VM-Series plugin operation commands. *Multiple commands must be entered on a single, comma separated with no spaces.* • **sriov-access-mode-on**—this command is only valid for VM-Series firewall on ESXi and KVM hypervisors. For KVM only, if you enable **sriov-access-mode-on**, do not enable **op-command-modes=jumbo-frame**. • **aws-gwlb-inspect:enable**—enables VM-Series Integration with an AWS Gateway Load Balancer. • **aws-gwlb-associate-vpce:<vpce-id>@ethernet<subinterface>**—allows you to Associate a VPC Endpoint with a VM-Series Interface or subinterface on the firewall. The specified interface is assigned to a security zone. • **aws-gwlb-overlay-routing:enable**—use this command to Enable Overlay Routing for the VM-Series on AWS when integrated with a AWS GWLB. • **set-dp-cores—customize the number of cores for a VM with a Software NGFW license. The maximum number of cores is 16.**

 ©2024 Palo Alto Networks, Inc.

Field	Description
dhcp-send-hostname=	The value of yes or no comes from the DHCP server. If yes, the firewall will send its hostname to the DHCP server. This field is relevant only if type is dhcp-client.
dhcp-send-client-id=	The value of yes or no comes from the DHCP server. If yes, the firewall will send its client ID to the DHCP server. This field is relevant only if type is dhcp-client.
dhcp-accept-server-hostname=	The value of yes or no comes from the DHCP server. If yes, the firewall will accept its hostname from the DHCP server. This field is relevant only if type is dhcp-client.
dhcp-accept-server-domain=	The value of yes or no comes from the DHCP server. If yes, the firewall will accept its DNS server from the DHCP server. This field is relevant only if type is dhcp-client.
vm-series-auto-registration-pin-id and vm-series-auto-registration-pin-value	The VM-Series registration PIN ID and value for installing the device certificate on the VM-Series firewall. The PIN ID and value also enable you to automatically activate the site licenses for AutoFocus and Cortex Data Lake on PAYG instances of the firewall. You must generate this in registration PIN ID and value on the Palo Alto Networks CSP. See Install a Device Certificate on the VM-Series Firewall for information on generating PIN ID and value.

Sample init-cfg.txt File

The following sample basic configuration files show all the parameters that are supported in the file; required parameters are in **bold**.

Sample init-cfg.txt file (Static IP Address)	Sample init-cfg.txt file (DHCP Client)
type=static	**type=dhcp-client**
ip-address=10.*.*.19	ip-address=
default-gateway=10.*.*.1	default-gateway=
netmask=255.255.255.0	netmask=
ipv6-address=2001:400:f00::1/64	ipv6-address=
ipv6-default-gateway=2001:400:f00::2*	ipv6-default-gateway=
hostname=Ca-FW-DC1	hostname=Ca-FW-DC1
vm-auth-key=7550362253****	vm-auth-key=7550362253****
panorama-server=10.*.*.20	panorama-server=10.*.*.20
panorama-server-2=10.*.*.21	panorama-server-2=10.*.*.21

 ©2024 Palo Alto Networks, Inc.

Sample init-cfg.txt file (Static IP Address)	Sample init-cfg.txt file (DHCP Client)
tplname=FINANCE_TG4	tplname=FINANCE_TG4
dgname=finance_dg	dgname=finance_dg
dns-primary=10.5.6.6	dns-primary=10.5.6.6
dns-secondary=10.5.6.7	dns-secondary=10.5.6.7
op-command-modes=jumbo-frame,mgmt-interface-swap**	op-command-modes=jumbo-frame,mgmt-interface-swap**
op-cmd-dpdk-pkt-io=***	op-cmd-dpdk-pkt-io=***
plugin-op-commands=	plugin-op-commands=
dhcp-send-hostname=no	dhcp-send-hostname=yes
dhcp-send-client-id=no	dhcp-send-client-id=yes
dhcp-accept-server-hostname=no	dhcp-accept-server-hostname=yes
dhcp-accept-server-domain=no	dhcp-accept-server-domain=yes
vm-series-auto-registration-pin-id=abcdefgh1234**	**vm-series-auto-registration-pin-id=abcdefgh1234****
vm-series-auto-registration-pin-value=zyxwvut-0987**	**vm-series-auto-registration-pin-value=zyxwvut-0987****

You cannot specify the management IP address and netmask configuration for the VM-Series firewall on AWS. If defined, the firewall ignores the values you specify because AWS uses a back-end metadata file to assign the management IP address and netmask.

**The IPv6 default gateway is required if you include an IPv6 address.*

***The `mgmt-interface-swap` operational command pertains only to a VM-Series firewall on AWS or GCP.*

****The `op-cmd-dpdk-pkt-io=off` is for disabling DPDK on the VM-Series firewall on ESXi, KVM, and GCP (DPDK is enabled by default).*

***** The `vm-series-auto-registration-pin-id` and `vm-series-auto-registration-pin-value` are required for two use cases:*

- *Activation of site licenses—AutoFocus or Cortex Data Lake—with Pay-As-You-Go (PAYG) license options of the VM-Series firewall.*
- *Retrieve and install the device certificate on the VM-Series firewall.*

 ©2024 Palo Alto Networks, Inc.

Create the bootstrap.xml File

Use these instructions to export the configuration from a firewall running on the same platform or hypervisor as your target deployment.

STEP 1 | Export a configuration from a firewall.

1. Select **Device** > **Setup** > **Operations**.
2. Select the configuration file you want to export.
   - To export the running configuration, in the Configuration Management section, **Export named configuration snapshot** and select **running config.xml** from the drop-down.
   - To export a previous version of a firewall configuration, in the Configuration Management section, **Export configuration version** and select the appropriate configuration version in the drop-down.

STEP 2 | Rename the configuration file and save.

1. Rename the file to `bootstrap.xml`.

   For the bootstrap process to be successful, the filename must be an exact (case-sensitive) match.
2. Save the `bootstrap.xml` file in the same location as the `init-cfg.txt` file.

 ©2024 Palo Alto Networks, Inc.

Prepare the Licenses for Bootstrapping

To license the firewall during the bootstrapping process, you must purchase the auth codes and register the licenses and subscriptions on the Palo Alto Networks Support portal before you begin bootstrapping.

For the VM-Series firewalls running BYOL (not applicable for usage-based licensing—PAYG), you must have an auth code bundle that includes the capacity auth code, support subscription, and any other subscriptions you require. The process of preparing the licenses for bootstrapping depends on whether the firewall has Internet access when bootstrapping:

- Direct Internet access—The firewall is connected directly to the Internet.
- Indirect Internet access—The firewall is managed by Panorama, which has direct Internet access and the ability to fetch the license keys on behalf of the firewall.
- No Internet access—The firewall uses an orchestration service or a custom script to fetch the license keys on behalf of the firewall.

- For VM-Series firewalls with Internet access.

  Enter the auth code in the `/license` folder when you Prepare the Bootstrap Package.

- For VM-Series firewalls with indirect Internet access.
  1. Register the auth code on the Palo Alto Networks Support portal.
     1. Go to the Support portal, log in, and select **Assets** > **Register New Device** > **Register device using Serial Number or Authorization Code**.
     2. Follow the steps to Register the VM-Series Firewall.
     3. Click **Submit**.
  2. Activate the auth codes on the Palo Alto Networks Support portal to generate license keys.
     1. Go to the Support portal, log in, and select the **Assets** tab.
     2. For each serial number, click the **Action** link.
     3. Select the **Activate Auth-Code** button.
     4. Enter the **Authorization code**, click **Agree**, and **Submit**.
     5. Download the license keys and save it to a local folder.
     6. Continue to Prepare the Bootstrap Package; you must add the license keys that you downloaded to the `\license` folder in the bootstrap package.

- For a custom script or an orchestration service that can access the Internet on behalf of firewalls.

  The script or service must fetch the CPU ID and the UUID from the hypervisor on which the firewall is deployed and access the Palo Alto Networks Support portal with CPU ID, UUID, API key and the auth code to obtain the required keys. See Licensing API.

 ©2024 Palo Alto Networks, Inc.

Prepare the Bootstrap Package

On AWS, Azure, or GCP, you can create the bootstrap package in your public cloud storage.

- VM-Series plugin version 1.0.13 and earlier, and versions 2.0.0 and 2.0.1 support one bootstrap package per storage bucket.
- VM-Series plugin versions 2.0.2 and later also support subfolders within your public cloud storage bucket. Within a bucket you can create multiple folders and subfolders, each containing a bootstrap package. Typically a folder represents configuration for a group of VMs, such as a Panorama device group.

  To access the bootstrap package, specify the full path to the bootstrap folder. For example: `my-storage/my-firewalls/bootstrap-2020-10-15`

Use the following procedure to prepare the bootstrap package.

STEP 1 | Create the top-level directory structure for the bootstrap package.

On your local client or laptop, or in a public cloud storage bucket, create the following folders:

```
/config
/content
/software
/license
/plugins
```

You can leave a folder empty, but you must have **/config**, **/license**, **/software**, and **/content** folders. The **/plugins** folder is optional, and only required if you are upgrading the VM-Series plugin independent of a PAN-OS release.

Do not place any other files or folders in the bootstrap structure. Adding other files or folders will result in a bootstrapping failure.

```
/my-storage
   /my-firewalls
      /internal     /external
        /config        /config
        /content       /content
        /license       /license
        /plugins       /plugins
        /software      /software
```

STEP 2 | Add content within each folder.

For an overview of the process, see Bootstrap Package. For details on the files in the **/config** folder, see Bootstrap Configuration Files.

```
/config
  0008C100105-init-cfg.txt
  0008C100107-init-cfg.txt
  bootstrap.xml
```

 ©2024 Palo Alto Networks, Inc.

```
/content
  panupv2-all-contents-488-2590
  panup-all-antivirus-1494-1969
  panup-all-wildfire-54746-61460
/software
  PanOS_vm-10.0.0
/license
  authcodes
  0001A100110-url3.key
  0001A100110-threats.key
  0001A100110-url3-wildfire.key
/plugins
  vm_series-2.0.2
```

- *If you save the keys to the `license` folder, you can use a file naming convention that works for you, but keep the `.key` extension in the filename. For auth codes, create a text file named `authcodes` (without a file extension), add your auth codes to that file, and save it to the `license` folder.*
- *Use an auth code bundle instead of individual auth codes so that the firewall or orchestration service can simultaneously fetch all license keys associated with a firewall. If you use individual auth codes instead of a bundle, the firewall will retrieve only the license key for the first auth code included in the file.*
- *In the `/plugins` folder, supply only one VM-Series plugin binary. Do not supply multiple plugin versions.*

STEP 3 | Create the bootstrap package.

For VM-Series firewalls, create the image in the appropriate format for your hypervisor. See Bootstrap Package Delivery.

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on AWS

STEP 1 | Choose a bootstrap method.

- To add a basic configuration to the bootstrap package, continue to Step 2.
- To enter a basic configuration as user data or use user data to get the basic configuration from an AWS secret, continue to Step 3.

STEP 2 | Prepare an S3 bucket, and an IAM role to enable read access.

To bootstrap using a file, you must be familiar with AWS S3 and IAM permissions required for completing this process. For detailed instructions on creating policy, refer to the AWS documentation on Creating Customer Managed Polices.

The management interface of the VM-Series firewall must be able to access the S3 bucket to complete bootstrapping. You can either assign a public IP address or an elastic IP address to the management interface so that the S3 bucket can be accessed over the Internet. Or, create a AWS VPC endpoint in the same region as the S3 bucket, if you prefer to create a private connection between your VPC and the S3 bucket and do not want to enable internet access on the firewall management interface. For more information refer to the AWS documentation on setting up VPC endpoints.

If PA-VM is deployed in any of the non-default regions such as Asia Pacific (Hong Kong) and Africa (Cape Town), bootstrapping works only if you use local S3 buckets. This is inline with the default AWS behavior for those regions to restrict interactions with the AWS resources on other regions.

1. Create an IAM role with inline policy to enable read access to the S3 bucket [ListBucket, GetObject]. For detailed instructions on creating an IAM role, defining which accounts or AWS services can assume the role, defining which API actions and resources the application can use upon assuming the role, refer to the AWS documentation on IAM Roles for Amazon EC2. When launching the VM-Series firewall, you must attach this role to enable access to the S3 bucket and the objects included in the bucket for bootstrapping successfully.
2. On the AWS console, create an Amazon Simple Storage Service (S3) bucket, or create a subdirectory in an existing S3 bucket.

   The S3 bucket in the following example, vmseries-aws-bucket, is at the All Buckets root folder level.

```
{
   "Version": "2012-10-17",
   "Statement": [
   {
      "Effect": "Allow",
      "Action": ["s3:ListBucket"],
      "Resource": ["arn:aws:s3:::<bucketname>"]
   },
   {
      "Effect": "Allow",
      "Action": ["s3:GetObject"],
```

 ©2024 Palo Alto Networks, Inc.

```
        "Resource": ["arn:aws:s3:::<bucketname>/*"]
        }
    ]
}
```

3. Create the folders within the S3 bucket as described in Prepare the Bootstrap Package.
   - Create the structure directly in your S3 bucket.

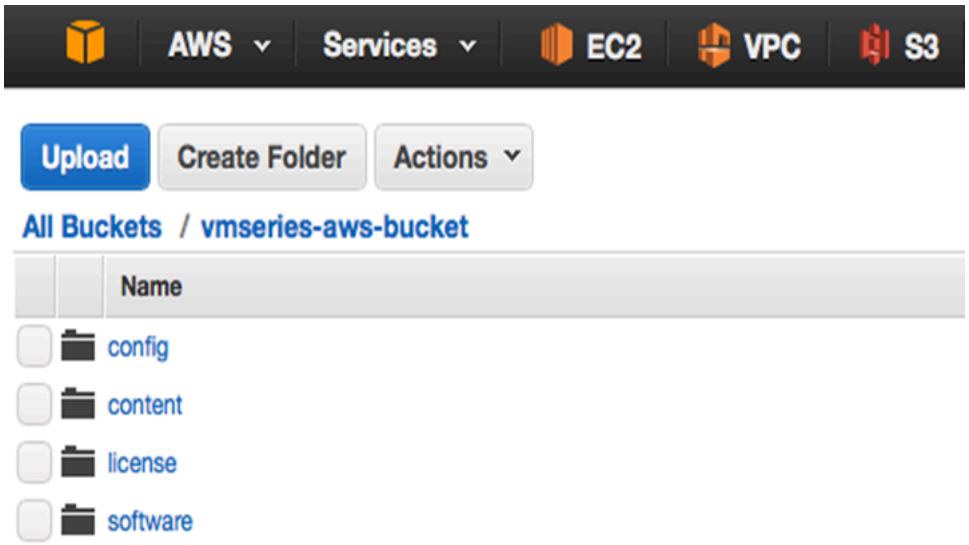

   - (Optional) Add content within each folder. You can leave a folder empty, but you must have all the \config, \content, \license, and \software folders. The \plugins folder is optional.

> *If you have enabled logging in Amazon S3, a Logs folder is automatically created in the S3 bucket. The Logs folder helps troubleshoot issues with access to the S3 bucket.*

STEP 3 | Launch the VM-Series firewall on AWS. Choose one of the following.

- **`init-cfg.txt`**—If you are using a file to configure the firewall, attach the IAM role you created in Step 2.1, expand the **Advanced Details** section, and in the **User Data** field, specify the path to an S3 bucket, directory, or subdirectory. For example,

```
vmseries-bootstrap-aws-s3bucket=<bucketname>
```

  or

```
vmseries-bootstrap-aws-s3bucket=<bucketname/directoryname>
```

- **User Data**—If you are using user data to configure the firewall, expand the **Advanced Details** section and in the **User Data** field enter the initial bootstrap parameters as described in Enter a Basic Configuration as User Data (Public Clouds).
- **AWS Secrets Manager**—If you stored your basic configuration as described in Save a Basic Configuration in the AWS Secrets Manager, expand the **Advanced Details** section and

 ©2024 Palo Alto Networks, Inc.

in the **User Data** field choose **As text** and enter the secret name as a key-value pair. For example:

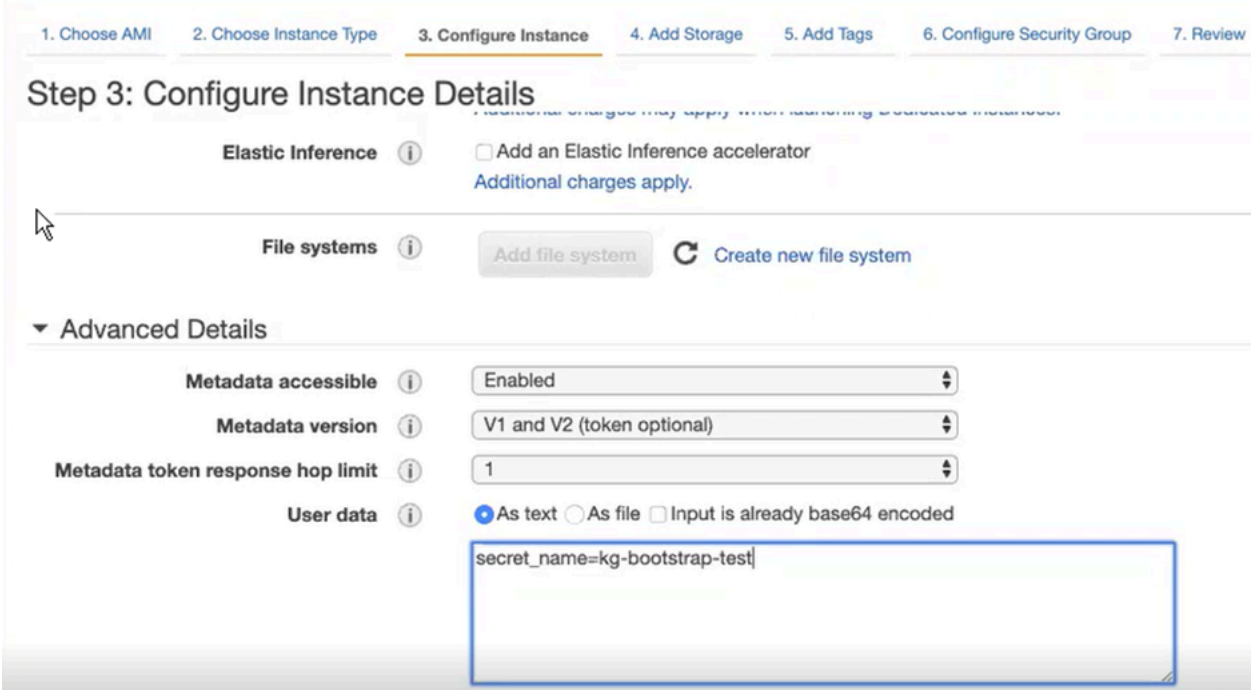

Select **Review and Launch**. For more details, see Launch the VM-Series firewall on AWS.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Verify Bootstrap Completion. Select the firewall instance on the AWS Management console and choose **Actions** > **Instance Settings** > **Get Instance Screenshot**.

- The screenshot shows bootstrapping in progress. A successful bootstrap is shown below:

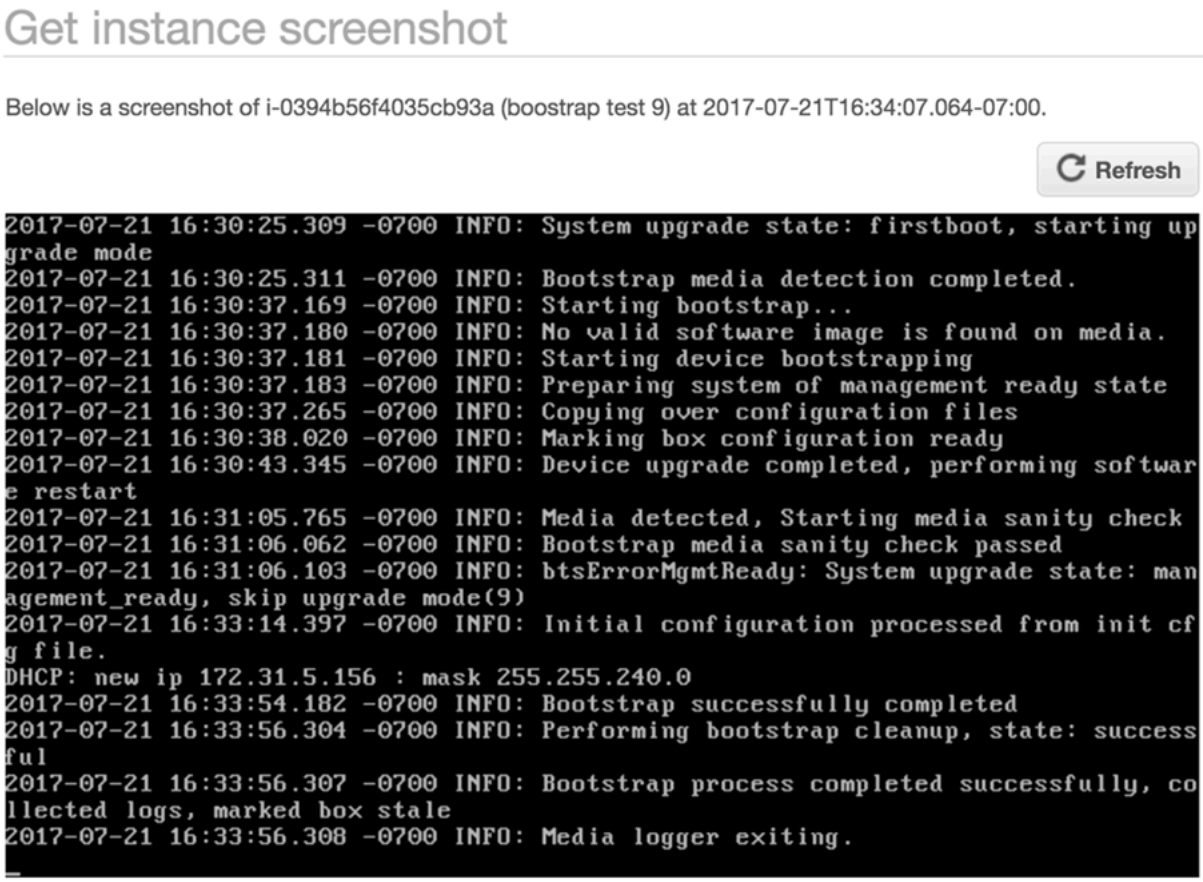

- If you are using an S3 bucket and the S3 bucket does not have the correct permissions or you do not have all four folders in the S3 bucket, you see the following error message:

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on Azure

The VM-Series firewall on Azure supports *Azure Files* service for bootstrapping.

STEP 1 | Choose a bootstrap method.

- To add a basic configuration to the bootstrap package, continue to Step 2.

  To manage the bootstrap package for the VM-Series firewall on Azure, you must be familiar with storage accounts on Azure and know how to create a file share and directory objects that contain the folder structure required for the bootstrap package. You can share an Azure file share across many virtual machines so that all firewalls deployed in the same region as the storage account that hosts the file share can access the files concurrently.

  The management interface of the VM-Series firewall must be able to access the file share that holds the bootstrap package so that it can complete bootstrapping.

- To Enter a Basic Configuration as User Data (Public Clouds), continue to Step 3.2.

STEP 2 | Set up the bootstrap package within an Azure Files service.

1. On the Azure portal, select or create a storage account.
2. Create a file share within the Azure Files service.

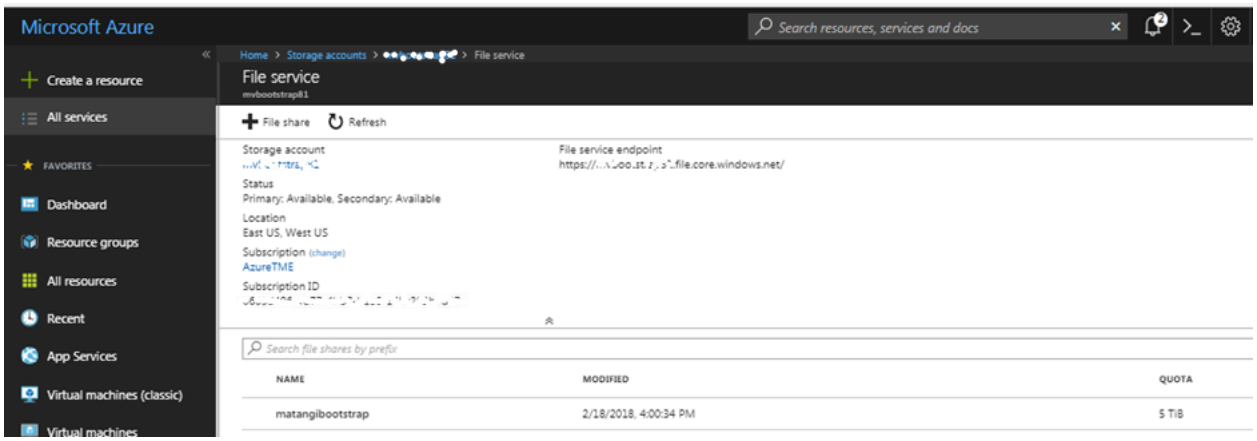

3. Create the folders within the storage account.
   - Create the top-level directory structure for the bootstrap package directly in the root folder and create a subfolder for each bootstrap configuration.
   - Add content folders within each folder. You can leave a folder empty but you must have all four folders (config, license, software and content) in the parent folder. In the

 ©2024 Palo Alto Networks, Inc.

following screenshot, you can see that the `config` folder has the `init-cfg.txt` file uploaded to it.

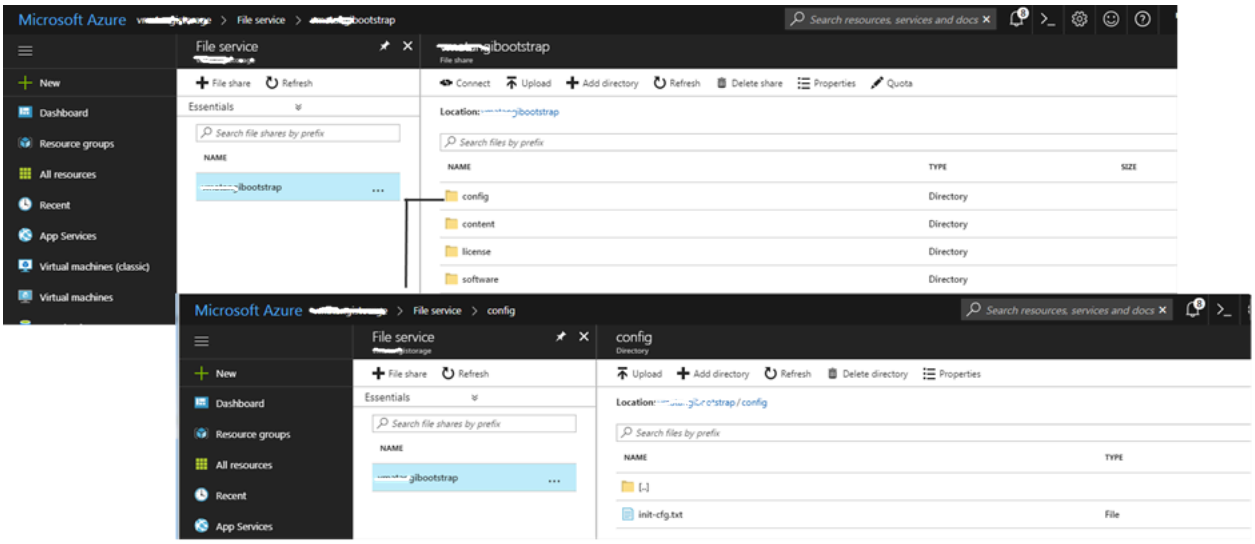

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Deploy the VM-Series Firewall from the Azure Marketplace (Solution Template).

- If you are using a file to configure the firewall, continue to Step 3.1
- If you are using custom data to configure the firewall, continue to Step 3.2.

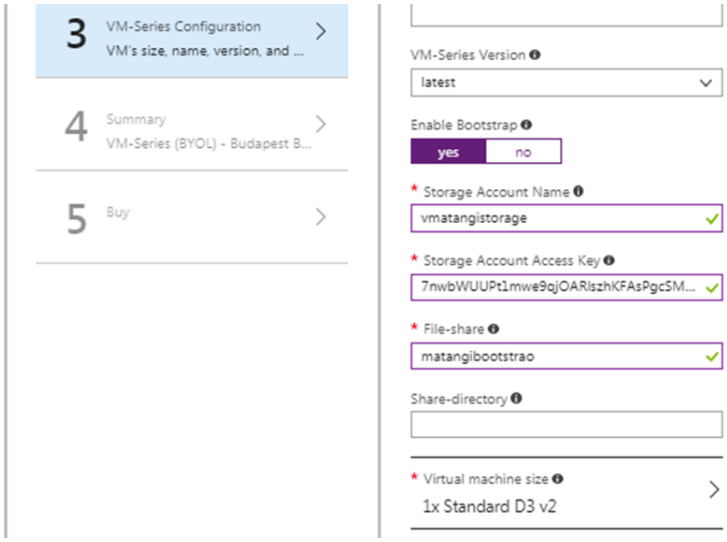

1. If you choose to use the bootstrap package, select **Enable Bootstrap: Yes** and provide the information required to access the file share that holds the bootstrap files.
    1. **Storage Account Name**— This is the Azure storage account in which you created the file share for the bootstrap folders.
    2. **Storage Account Access Key**—The firewall needs this access key to authenticate to the storage account and access the files stored within. To copy this access key, select the storage account name, and then select **Settings** > **Access Keys**.

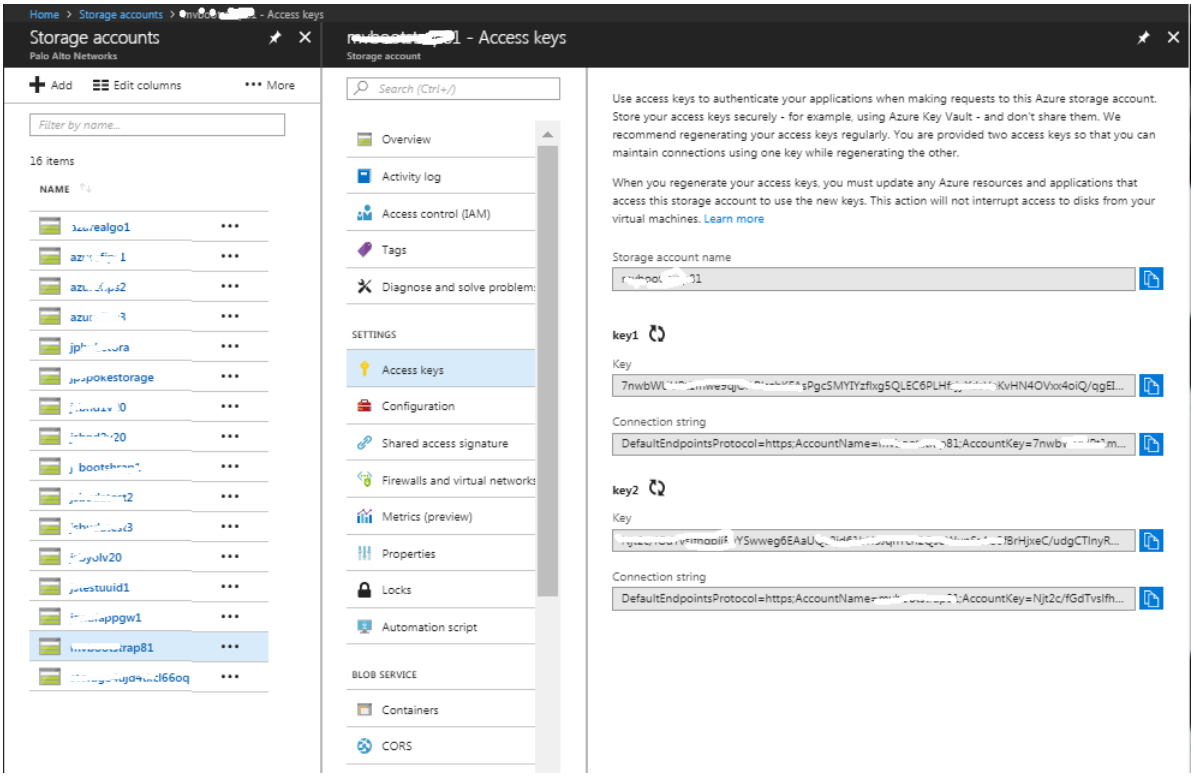

 ©2024 Palo Alto Networks, Inc.

3. **File-share**—The file-share name that contains the bootstrap package.
4. (Optional) **Share-directory**—The path to a subfolder within the file-share. If you have a common file share that serves as a repository for bootstrap configurations for different set ups, you can use a share-directory to create a folder hierarchy and access a specific set of subfolders within the common file-share.

2. Enter the configuration parameters as custom data. For the key-value pairs, see Enter a Basic Configuration as User Data (Public Clouds). Separate each key-value pair with a semicolon. For example:

```
type=dhcp-client; op-command-modes=jumbo-frame;
vm-series-auto-registration-pin-id=abcdefgh1234****;
vm-series-auto-registration-pin-value=zyxwvut-0987****
```

Provide custom data using one of the methods in Custom data and Cloud-Init on Azure Virtual Machines.

STEP 4 | Verify Bootstrap Completion.

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on ESXi

You can bootstrap the VM-Series firewall using an ISO image or a virtual hard disk.

- Bootstrap the VM-Series Firewall on ESXi with an ISO
- Bootstrap the VM-Series Firewall on ESXi with a Block Storage Device

Bootstrap the VM-Series Firewall on ESXi with an ISO

Use these instructions to bootstrap the VM-Series firewall on an ESXi server using an ISO.

STEP 1 | Create an ISO image and upload it to a Virtual Machine File System (VMFS) datastore or to a Network File System (NFS) volume.

1. Prepare the Bootstrap Package.
2. Create an ISO image. The tool you use to create the image varies based on your client operating system.
3. Upload the ISO image to a VMFS datastore or to an NFS volume that is accessible to the ESX/ESXI host.

STEP 2 | Deploy the firewall.

1. Provision the VM-Series Firewall on an ESXi Server.

   By default, the firewall is deployed with two network interfaces— one for management traffic and one data traffic. Make sure that the first ethernet interface on the firewall, which is its management interface, is connected to the virtual switch port-group assigned for device management.
2. Do not power on the firewall.

STEP 3 | Attach the bootstrap image to the firewall.

1. Select the VM-Series firewall from the **Inventory** list.
2. Click **Edit Settings** and select **Virtual Hardware**.
3. Select **Datastore iso file** in the **CD DVD drive** drop-down, and **browse** for the ISO image.
4. Power on the firewall. The firewall will begin with the bootstrapping process, which will take several minutes. The status messages on the success or failure of the process will display on the console.
5. Verify Bootstrap Completion.

Bootstrap the VM-Series Firewall on ESXi with a Block Storage Device

Use these instructions to bootstrap the VM-Series firewall on an ESXi server using a block storage device.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create the bootstrap package and the virtual hard disk.

1. Create the bootstrap package.
2. Deploy a Linux virtual machine.
3. On the Linux machine, Prepare the Bootstrap Package. You can leave the folder empty, but you must have all four folders.
4. Attach a new data disk less than 39 GB to the Linux virtual machine.
5. Partition the disk and format the file system as ext3.
6. Make a directory for the new file system and mount the disk to the Linux virtual machine.
7. Copy the contents of your bootstrap package to the disk.
8. Unmount the disk.
9. Detach the disk from the Linux virtual machine. Take note of the Disk File describing the bootstrap disk you created; it shows the datastore name and path to the disk. Additionally, do not check the Delete Files From Datastore check box; doing so deletes the disk.

STEP 2 | Deploy the firewall.

1. Provision the VM-Series Firewall on an ESXi Server.
2. Do not power on the firewall.

STEP 3 | Attach the bootstrap package to the firewall.

1. Select the VM-Series firewall from the Inventory list.
2. Click **Edit Settings** and select **Virtual Hardware**.
3. From the New Device drop-down, select **Existing Hard Disk**. Select the bootstrap disk according to the datastore and path noted previously.
4. Power on the firewall. The firewall will begin with the bootstrapping process, which will take several minutes. The status messages on the success or failure of the process will display on the console.
5. Verify Bootstrap Completion.

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on Google Cloud Platform

STEP 1 | Choose a bootstrap method.

STEP 2 | Log in to Google Cloud Console.

- To add a basic configuration to the bootstrap package, create bootstrap files as described in Prepare the Bootstrap Package, and continue to Step 3.
- To enter a basic configuration as custom metadata, skip to Step 4.

STEP 3 | Select **Storage** > **Browser**, and click **Create Bucket**.

You can use this bucket to bootstrap the firewall when you Deploy the VM-Series Firewall from Google Cloud Platform Marketplace.

If you intend to bootstrap using a Google storage bucket in the same project, you must have devstorage.read_only IAM privileges.

You can create and populate the bucket at the top level, or you can create a bucket with subfolders within it so that many bootstrap packages can share the same bucket.

1. Enter the bucket name, choose the default storage class, and choose a location. Note, the location in the storage bucket must be compatible with the zone you specify for the compute engine instance.
2. Click **Create**.
3. In the Storage Browser, click the bucket name to open it.
4. Click **Create Folder** and name the folder `config`. Click **Create**.
5. Repeat, creating folders for `content`, `license`, and `software`, as shown below. All folders must be present, even if they are empty.

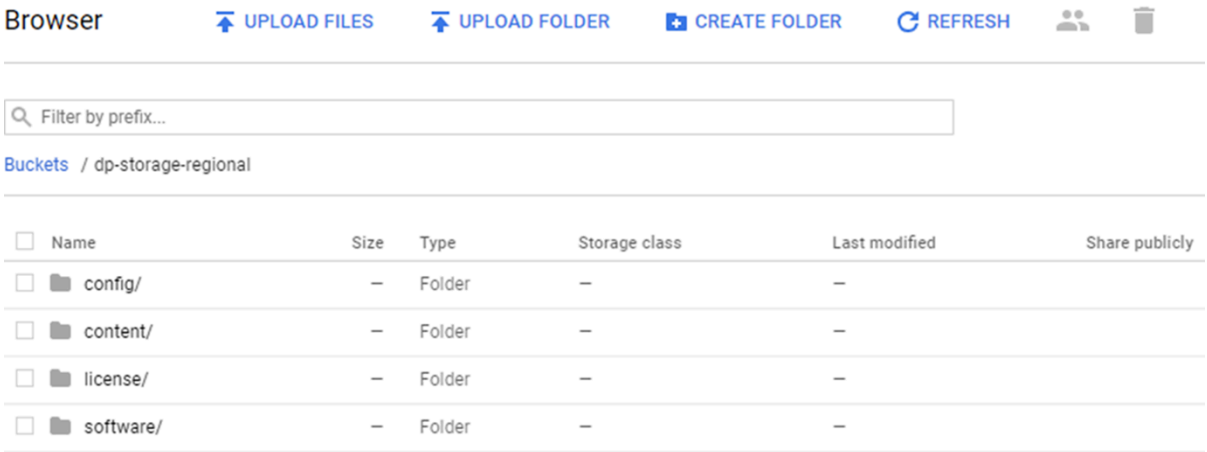

6. (Optional) If you created an **`init-cfg.txt`** file, open the `config` folder. Click **Upload Files**, browse to select your **`init-cfg.txt`** file, and click **Open**.
7. Open the `license` folder and upload the **`authcodes`** file.
8. Continue until you have uploaded all the bootstrap files.

 ©2024 Palo Alto Networks, Inc.

STEP 4 | Add the initial configuration parameters as metadata. **Add** each key-value pair as described in Enter a Basic Configuration as User Data (Public Clouds).

STEP 5 | See Deploy the VM-Series Firewall from Google Cloud Platform Marketplace for deployment details.

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on Hyper-V

You can bootstrap the VM-Series firewall using an ISO image or a virtual hard disk.

- Bootstrap the VM-Series Firewall on Hyper-V with an ISO
- Bootstrap the VM-Series Firewall on Hyper-V with a Block Storage Device

Bootstrap the VM-Series Firewall on Hyper-V with an ISO

Use these instructions to bootstrap the VM-Series firewall on a Hyper-V server with an ISO.

STEP 1 | Create an ISO image.

1. Prepare the Bootstrap Package.
2. Create an ISO image. The tool you use to create the image varies based on your client operating system.
3. Upload the ISO image to a location accessible to the Hyper-V host.

STEP 2 | Deploy the firewall.

1. Provision the VM-Series Firewall on a Hyper-V host with Hyper-V Manager.

   By default, the firewall is deployed with two network interfaces— one for management traffic and one data traffic. Make sure that the first ethernet interface on the firewall, which is its management interface, is connected to the vSwitch assigned for device management.
2. Do not power on the firewall.

STEP 3 | Attach the bootstrap image to the firewall.

1. In Hyper-V Manager, select the VM-Series firewall from the **Virtual Machines** list.
2. Click **Settings** > **Hardware** > **IDE Controller** > **DVD Drive**.

   If you have more than one DVD drive, the ISO image must be applied to the first drive.
3. Under Media, click the **Image file** radio button.
4. Click **Browse** and select your uploaded ISO image.
5. Click **Apply** and **Ok** to exit the virtual machine settings.
6. Power on the firewall. The firewall will begin with the bootstrapping process, which will take several minutes. The status messages on the success or failure of the process will display on the console.
7. Verify Bootstrap Completion.

Bootstrap the VM-Series Firewall on Hyper-V with a Block Storage Device

Use these instructions to bootstrap the VM-Series firewall on a Hyper-V server with a block storage device.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create the bootstrap package and the virtual hard disk.

1. Deploy a Linux virtual machine.
2. On the Linux machine, Prepare the Bootstrap Package. You can leave the folder empty, but you must have all four folders.
3. Attach a new data disk less than 39 GB to the Linux virtual machine.
    1. Power of the Linux virtual machine.
    2. In Hyper-V, select the Linux virtual machine from the Virtual Machines list.
    3. Select **Settings** > **Hardware** > **IDE Controller**.
    4. Select **Hard Drive** and click **Add**.
    5. Select **Virtual Hard Disk** and click **New**.
    6. Follow the on-screen instructions to create a new VHD. Note the name and path of the new VHD.
    7. Click **Apply** then **OK** to exit the virtual machine settings.
    8. Power on the Linux virtual machine.
4. Connect to the CLI of the Linux virtual machine.
5. Partition the disk and format the file system as ext3.
6. Make a directory for the new file system and mount the disk to the Linux virtual machine.
7. Copy the contents of your bootstrap package to the disk.
8. Unmount the disk.
9. Detach the disk from the Linux virtual machine.
    1. Power of the Linux virtual machine.
    2. Select the Linux virtual machine from the Virtual Machines list.
    3. Select **Settings** > **Hardware** > **IDE Controller**.
    4. Select the VHD you created.
    5. Click **Remove**. This detaches the VHD but does not delete it.

STEP 2 | Deploy the firewall.

1. Provision the VM-Series Firewall on a Hyper-V host with Hyper-V Manager.
2. Do not power on the firewall.

 ©2024 Palo Alto Networks, Inc.

STEP 3 | Attach the bootstrap disk image to the firewall.

1. Select the firewall from the Virtual Machines list.
2. Select **Settings** > **Hardware** > **IDE Controller**.
3. Select **Hard Drive** and click **Add**.
4. Select **Virtual Hard Disk** and click **Browse**.
5. Browse to the bootstrap VHD you created, select it, and click **Open**.
6. Click **Apply** and **OK** to exit the Virtual Machine settings.
7. Power on the firewall. The firewall will begin with the bootstrapping process, which will take several minutes. The status messages on the success or failure of the process will display on the console.
8. Verify Bootstrap Completion.

 ©2024 Palo Alto Networks, Inc.

Bootstrap the VM-Series Firewall on KVM

You can bootstrap the VM-Series firewall on KVM using an ISO image or a virtual hard disk.

- Bootstrap the VM-Series Firewall on KVM with an ISO
- Bootstrap the VM-Series Firewall on KVM With a Block Storage Device

Bootstrap the VM-Series Firewall on KVM with an ISO

Use these instructions to bootstrap the VM-Series firewall on a KVM server using an ISO.

STEP 1 | Create an ISO image.

1. Prepare the Bootstrap Package.
2. Create an ISO image. The tool you use to create the image varies based on your client operating system.
3. Upload the ISO image to a location accessible to the KVM host.

STEP 2 | Deploy the firewall.

1. Install the VM-Series Firewall on KVM.

   By default, the firewall is deployed with two network interfaces— one for management traffic and one data traffic. Make sure that the first ethernet interface on the firewall, which is its management interface, is connected to the virtual switch port-group assigned for device management.
2. Do not power on the firewall.

STEP 3 | Attach the bootstrap image to the firewall.

1. In virt-manager, double-click on the VM-Series firewall to open the console.
2. View the VM hardware details by navigating to **View** > **Details**.
3. Open the Add New Virtual Hardware menu by clicking **Add Hardware**.
4. Change the device type to IDE CDROM.
5. Click the **Select managed or other existing storage** radio button and click **Browse**. Locate the ISO image you created and click **Choose Volume**.
6. Click **Finish** to exit the Add New Virtual Hardware menu.
7. Power on the firewall by navigating to **Virtual Machine** > **Run**. The firewall will begin with the bootstrapping process, which will take several minutes. The status messages on the success or failure of the process will display on the console.
8. Verify Bootstrap Completion.

Bootstrap the VM-Series Firewall on KVM With a Block Storage Device

Use these instructions to bootstrap the VM-Series firewall on a KVM server with a block storage device.

 ©2024 Palo Alto Networks, Inc.

STEP 1 | Create the bootstrap package and the virtual hard disk.

1. Create the bootstrap package.
2. Create a new disk image less than 39 GB in size and partition the disk and format the file system as ext3. The tools used to complete this process vary based on your client operating system.
3. Mount the disk image file and copy the prepared bootstrap package to the disk image files.
4. Copy the contents of your bootstrap package to the disk.
5. Unmount the disk image.
6. Upload the disk image file to a location accessible to the KVM host.

STEP 2 | Deploy the firewall.

1. Install the VM-Series Firewall on KVM.
2. Do not power on the firewall.

STEP 3 | Attach the bootstrap disk image to the firewall.

1. In virt-manager, double click on the VM-Series firewall to open the console.
2. View the VM hardware details be selecting **View** > **Details**.
3. Open the Add New Virtual Hardware menu by clicking **Add Hardware**.
4. Select **Storage** and the select **Select or create custom storage**.
5. Click the **Manage** button to open the **Choose Storage Volume** dialog, and select the disk image file that you previously created.
6. Click Choose Volume.
7. Ensure that the device type is Disk Device and do not change the Bus Type.
8. Click Finish.
9. Power on the firewall. The firewall will begin with the bootstrapping process, which will take several minutes. The status messages on the success or failure of the process will display on the console.
10. Verify Bootstrap Completion.

 ©2024 Palo Alto Networks, Inc.

Verify Bootstrap Completion

You can see basic status logs on the console during the bootstrap and you can verify that the process is complete.

STEP 1 | If you included **`panorama-server`**, **`tplname`**, and **`dgname`** in your `init-cfg.txt` file, check Panorama managed devices, the device group, and the template name.

STEP 2 | Verify the general system settings and configuration. Access the web interface and select **Dashboard** > **Widgets** > **System** or use the CLI operational commands **`show system info`** and **`showconfig running`**.

STEP 3 | Verify the license installation. Select **Device** > **Licenses** or use the CLI operational command **`request license info`**.

STEP 4 | If you have Panorama configured, manage the content versions and software versions from Panorama. If you do not have Panorama configured, use the web interface to manage content versions and software versions.

 ©2024 Palo Alto Networks, Inc.

Bootstrap Errors

If you receive an error message during the bootstrapping process, refer to the following table for details.

Error message (Severity)	Reasons
Boot image error (high)	• No external device was detected with the bootstrap package. Or • A critical error happened while booting from the image on the external device. The bootstrap process was aborted.
No bootstrap config file on external device (high)	The external device did not have the bootstrap configuration file.
Bad or no parameters for mandatory networking information in the bootstrap config file (high)	The networking parameters required for bootstrapping were either incorrect or missing. The error message lists the value—IP address, netmask, default gateway—that caused the bootstrap failure.
Failed to install license key for file <license-key-filename> (high)	The license key could not be applied. This error indicates that the license key used was invalid. The output includes the name of the license key that could not be applied.
Failed to install license key using authcode <authcode> (high)	The license auth code could not be applied. This error indicates that the license auth code used was invalid. The output includes the name of the authcode that could not be applied.
Failed content update commits (high)	The content updates were not successfully applied.
USB media prepared successfully using given bundle (informational)	The bootstrap image has been successfully complied on the USB flash device. <username>: Successfully prepared the USB using bundle <bundlename>
Successful bootstrap (informational)	The firewall was successfully provisioned with the bootstrap configuration file. The output includes the license keys installed and the filename of the bootstrap configuration. On the VM-Series firewalls only, the PAN-OS version and content update version are also displayed.

Read about the Bootstrap Package and how to Prepare the Bootstrap Package.

 ©2024 Palo Alto Networks, Inc.

 ©2024 Palo Alto Networks, Inc.