# Paloalto

# VM-Series High Availability on OCB Flex Engine Installation and Deployment Guide

30th December 2020

Version 1.3

## document control

| date | version no. | author | change/addition |
|---|---|---|---|
| 10th January 2020 | 1.0 | Ahmad Samak | Creation |
| 16 January 2020 | 1.1 | Ahmed Samaha | Modification and update |
| 25 November 2020 | 1.2 | Ahmad Samak | Modification and update (Deployment of Palo Alto KVM release 9.1.3) |
| 30 December 2020 | 1.3 | Ahmad Samak | Modification and update (Software Update of Palo Alto KVM release 9.1.3 to release 9.1.4) |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

table of contents

# 1 References

| Reference | Description | Link to document |
|---|---|---|
| [1] | OCB FE Help Center | https://docs.prod-cloud-ocb.orange-business.com/en-us/index.html |
| [2] | VM-Series Deployment Guide | https://www.paloaltonetworks.com/documentation/81/virtualization/virtualization |

# 2 Introduction

For new applications and service deployment, many organizations are moving to the public cloud. Instead of developing new applications and running them on their on-premises hardware, these organizations are increasingly using infrastructure hosted and maintained by remote vendors. These Infrastructure-as-a-Service (IaaS) environments, originally used by startups or niche purposes by enterprises, are increasingly being used for applications that provide business differentiation. Applications deployed in public cloud IaaS environments are becoming more prevalent because they offer several productivity and scale benefits to an organization.

### Purpose of This Guide

Although IaaS providers are responsible for ensuring the security and availability of their infrastructure, ultimately, organizations are still responsible for the security of the applications and data. This reference architecture describes how an organization can use the Palo Alto Networks® VM-Series firewalls running PAN-OS to bring visibility, control, and protection to your applications built in Orange Flex Engine.

This document provides architectural guidance for solution architects and engineers who are familiar with the next-generation firewall but not Orange Flex Engine. It links the technical aspects of the Orange FE and Palo Alto Networks solution together before exploring the technical design models of the architecture. Use this guide as a roadmap for architectural discussions between Palo Alto Networks and your organization

# 3 Public Cloud and On-Premises Differences

Organizations generally move to the public cloud with the goals of increasing scale and reducing time to deployment. Achieving these goals requires application architectures built specifically for the public cloud. Before you can architect for the public cloud, you must understand how it is different from traditional on-premises environments.

## 3.1 Scaling Methods

Traditionally, organizations scale on-premises deployments through the purchase of devices that have increased performance capacity. Scaling up an on-premises deployment in this method makes sense because the devices are typically purchased to last year's requirements and must be sized to satisfy the performance requirements during their lifetime.

Public cloud environments focus on scaling out the deployment instead of scaling up. This architectural difference stems primarily from the capability of public cloud environments to dynamically increase or decrease the number of resources you have allocated. In the public cloud, infrastructure used to satisfy performance requirements can have a lifetime in minutes instead of years. Instead of purchasing extra capacity for use at some time in the future, the dynamic nature of the public cloud allows you to allocate just the right amount of resources required to service the application.

What this means in practice is that to architect an application for the cloud, you need to distribute functionality, and each functional area should be built to scale out as necessary. Typically, this means a load balancer distributes traffic across a pool of identically configured resources. When changes occur in the application traffic, the number of resources you have allocated to the pool can be increased or decreased dynamically. This design method provides scale and resiliency. However, the application architecture must take into account that the resources are transient. For example, the application state should not be stored in the networking infrastructure or in the frontend application servers. Instead, store state information on the client or persistent storage services.

The ability to scale a cloud architecture extends not only to the capacity of an application but also capacity to deploy applications globally. Scaling an application to a new region in a traditional on-premises deployment requires significant investment and planning. Public cloud architectures are location-agnostic and can be deployed globally in a consistent amount of time.

## 3.2 Reduced Time to Deployment

To achieve the goals of a reduced time to deployment you have to have a development and deployment process that is repeatable and reacts to changes quickly. DevOps workflows are the primary method for implementing this process. DevOps workflows are highly dependent on the ability to automate, as much as possible, the process of deploying a resource or application. In practice, this means the cloud infrastructure, as well as the resources running on it, needs to be able to be bootstrapped, configured, updated, and destroyed programmatically. Compared to traditional on-premises deployments where devices deployment, configuration, and operation happen manually, automated workflows in a public cloud environment can significantly reduce time to deployment.

In fact, automation is so core to cloud design that many cloud application architectures deploy new capabilities through the automated build-out of new resources instead of updating the existing ones. This type of cloud architecture provides a number of benefits, not the least of which is the ability phase in the changes to a subset of the traffic as well as the ability to quickly roll back the changes by redirecting traffic from the new resources to the old.

## 3.3 Security Integration

VM-Series firewalls enable you to securely implement scalable cloud architectures and reduce time to deployment. Capabilities of VM-Series firewalls leveraged to achieve this include:

- Application visibility—VM-Series firewalls natively analyze all traffic in a single pass to determine the application, content, and user identity. The application, content, and user are used as core elements of your security policy and for visibility, reporting, and incident investigation.

- Prevent advanced attacks at the application level—Attacks, much like many applications, can use any port, rendering traditional prevention mechanisms ineffective. VM-Series firewalls allow you to use Threat Prevention and the WildFire™ cloud-based threat analysis service to apply application-specific threat prevention policies that block exploits, malware, and previously unknown threats from infecting your cloud.

- Consistent policy and management—Panorama™ network security management enables you to manage your VM-Series deployments across multiple cloud environments, along with your physical security appliances, thereby ensuring policy consistency and cohesiveness. Rich, centralized logging and reporting capabilities provide visibility into virtualized applications, users, and content.

- Automation features to reduce time to deployment—VM-Series firewalls include management features that enable you to integrate security into your public cloud development projects. You can use bootstrapping to automatically provision a firewall with a working configuration, complete with licenses and subscriptions, and then auto-register itself with Panorama. Firewall performance metrics and health information can be published to Azure Application Insights, so you can create automate actions based on performance and usage patterns. To automate policy updates when workloads change, a fully documented XML API and dynamic address groups allow VM-Series firewalls to consume external data in the form of tags that can drive policy updates dynamically. The result is that new applications and next-generation security can be deployed simultaneously in an automated manner.

# 4 License Options

You can license VM-Series firewalls on OCB FE with licenses purchased through regular Palo Alto Networks channels.

Bring your own license (BYOL) and VM-Series ELA—A license that you purchase from a partner, reseller, or directly from Palo Alto Networks. VM-Series firewalls support all capacity, support, and subscription licenses in BYOL.

When using your own licenses, you license VM-Series firewalls like a traditionally deployed appliance, and you must apply a license authorization code. After you apply the code to the device, the device registers with the Palo Alto Networks support portal and obtains information about its capacity and subscriptions. Subscription licenses include Threat Prevention, PAN-DB URL Filtering, AutoFocus™, GlobalProtect, and WildFire.

To accelerate firewall deployment, the VM-Series enterprise licensing agreement (ELA) provides a fixed price licensing option allowing unlimited deployment of VM-Series firewalls with BYOL. Palo Alto Networks offers licenses in one and three-year term agreements with no true-up at the end of the term.

The VM-Series ELA includes four components:

- Your choice of single VM-Series model that you can deploy as many times as you want and in as many virtual environments as you want. All of your VM-Series ELA deployments use a single license authorization code, which allows for easier automation and simplifies the deployment of firewalls.
- Threat Prevention, WildFire, GlobalProtect and PAN-DB Subscriptions for every VM-Series firewall deployed as part of the VM-Series ELA.
- Unlimited deployments of Panorama as a virtual appliance.
- Support that covers all the components deployed as part of the VM-Series ELA.

✓ Whichever licensing model you chose will be permanent. After you deploy them, VM-Series firewalls cannot switch between the PAYG and bring-your-own-license (BYOL) licensing models. Switching between licensing models requires deploying a new firewall and migrating the configuration. Migration between evaluation, a regular license, and ELA is possible because they are all part of the BYOL licensing model.

# 5 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.

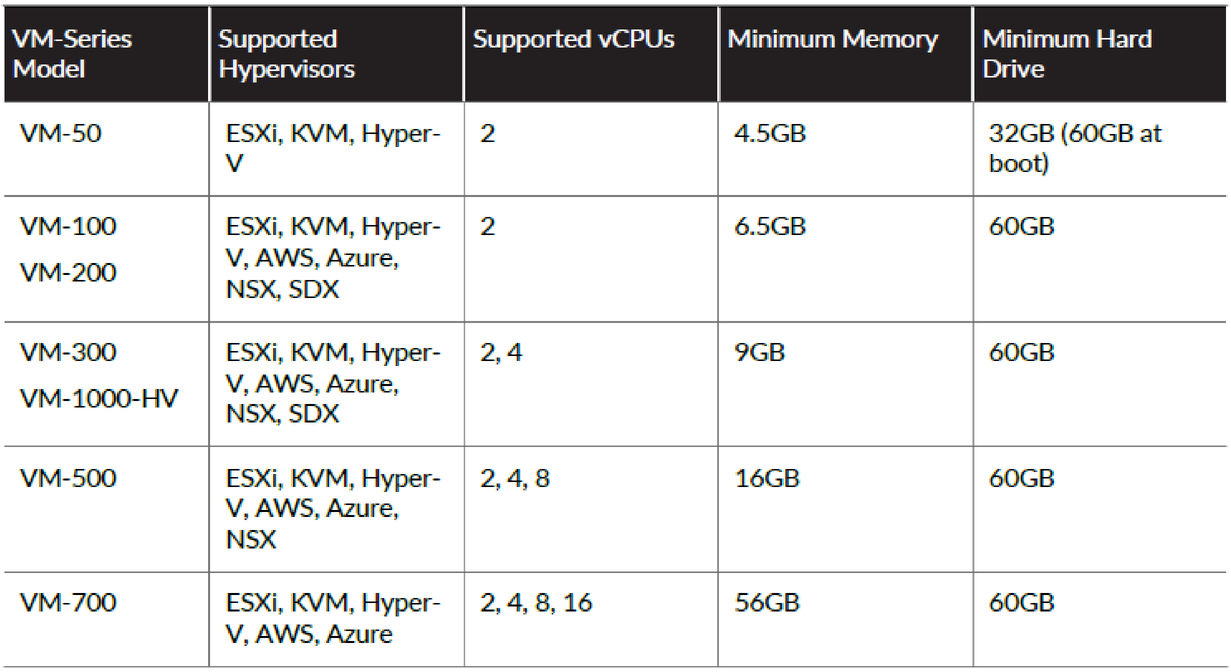

| VM-Series Model | Supported Hypervisors | Supported vCPUs | Minimum Memory | Minimum Hard Drive |
|---|---|---|---|---|
| VM-50 | ESXi, KVM, Hyper-V | 2 | 4.5GB | 32GB (60GB at boot) |
| VM-100<br>VM-200 | ESXi, KVM, Hyper-V, AWS, Azure, NSX, SDX | 2 | 6.5GB | 60GB |
| VM-300<br>VM-1000-HV | ESXi, KVM, Hyper-V, AWS, Azure, NSX, SDX | 2, 4 | 9GB | 60GB |
| VM-500 | ESXi, KVM, Hyper-V, AWS, Azure, NSX | 2, 4, 8 | 16GB | 60GB |
| VM-700 | ESXi, KVM, Hyper-V, AWS, Azure | 2, 4, 8, 16 | 56GB | 60GB |

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

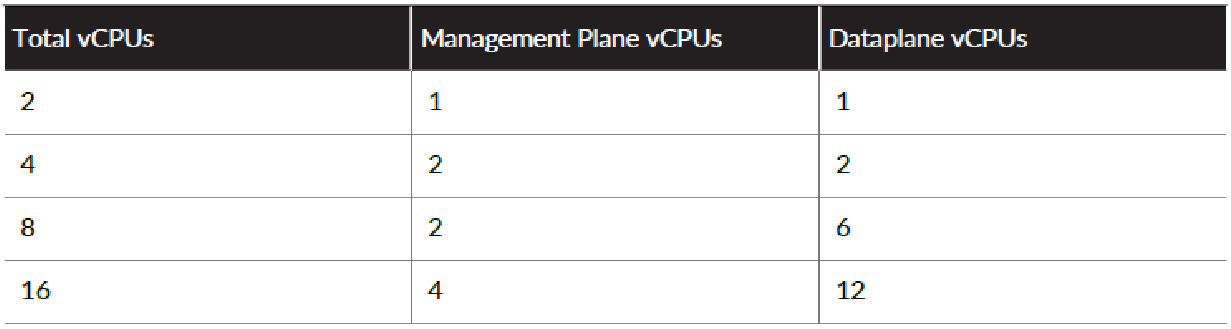

| Total vCPUs | Management Plane vCPUs | Dataplane vCPUs |
|---|---|---|
| 2 | 1 | 1 |
| 4 | 2 | 2 |
| 8 | 2 | 6 |
| 16 | 4 | 12 |

## 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-50sb 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.

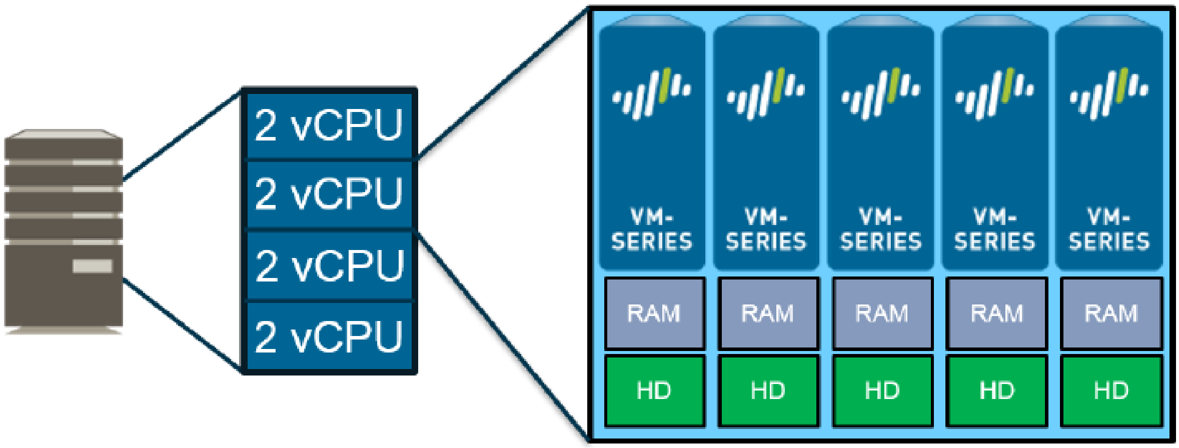

# 6 Deployment Method

Use the VM-Series firewall on OCB FE to secure your network users in the following scenarios:

## 6.1 Palo Alto High Availability (Active-Passive) model on OCB FE

The VM-Series firewall on OCB FE allows you to securely extend your physical data center/private cloud into OCB FE using IPsec tunneling. To improve your data center security, if you have segmented your network and deployed your workloads in separate VPC's, you can secure traffic flowing between VPC's with an IPsec tunnel and application whitelisting policies.

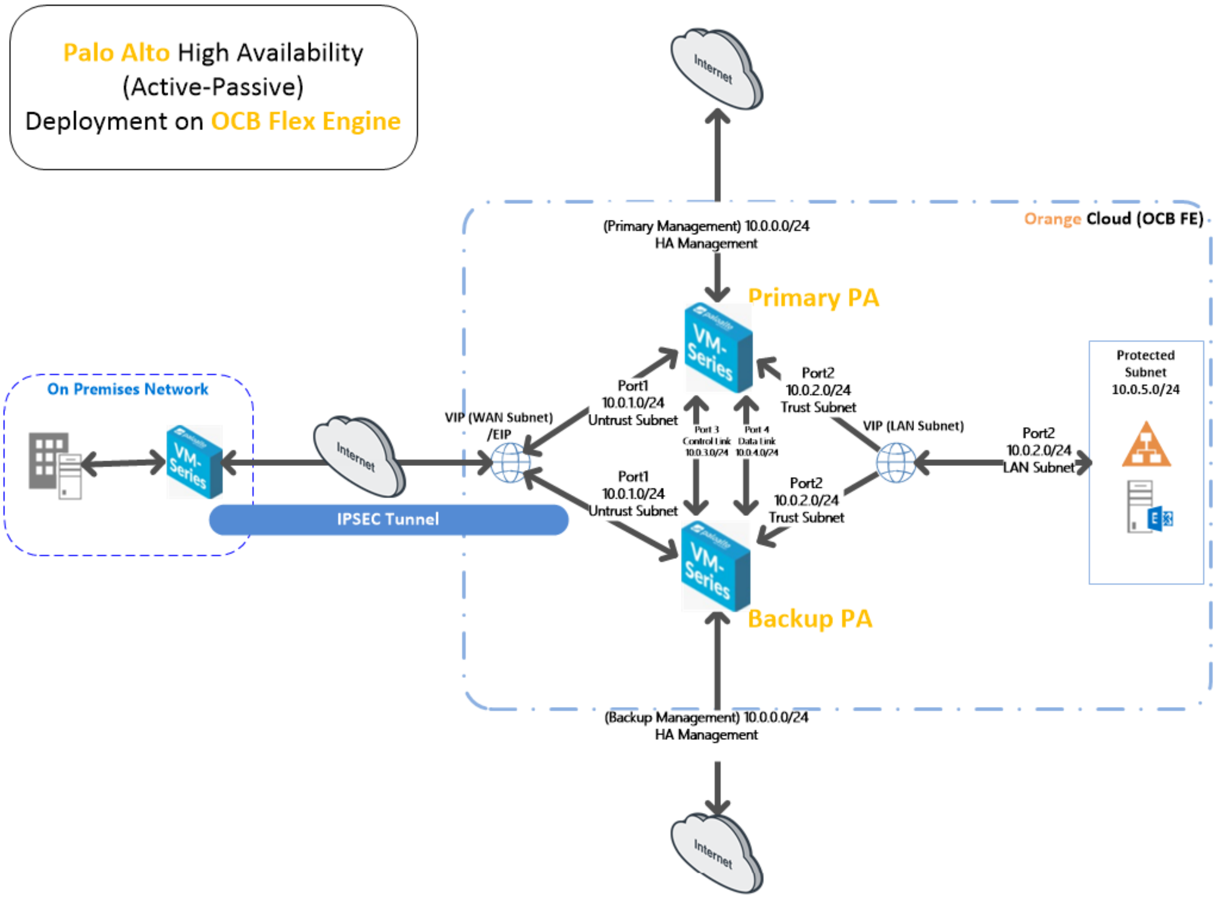

• Inter-Subnet —The VM-Series firewall can front your servers in a VPC and protects against lateral threats for inter-subnet traffic between applications in a multi-tier architecture.

• Gateway—The VM-Series firewall serves as the VPC gateway to protect Internet-facing deployments in the OCB FE (VPC). The VM-Series firewall secures traffic destined to the servers in the VPC and it also protects against lateral threats for inter-subnet traffic between applications in a multitier architecture.

• GlobalProtect—Use the OCB FE 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.

# 7 VM-Series High Availability on Orange Flex Engine

In our solution we have to deploy the following

- One VPC containing 6 or more Subnets
  - Management Subnet
  - Un-trust Subnet
  - Trust Subnet
  - Control link Subnet
  - Data Link Subnet
  - Protected Subnet
- Two Palo Alto VM-Series firewalls
- Two Virtual IP's (One in the un-trust Subnet and the other in the trust Subnet)

## 7.1 Create VPC

A VPC provides an isolated virtual network for ECSs. You can configure and manage the network as required.

To use a VPC, first create it by following the procedure provided in this section. Then, create subnets, security groups, and VPNs, and assign EIPs by following the procedure provided in subsequent sections based on your actual network requirements.

Procedure

1. Log in to the management console.
2. On the console homepage, under Network, click Virtual Private Cloud.

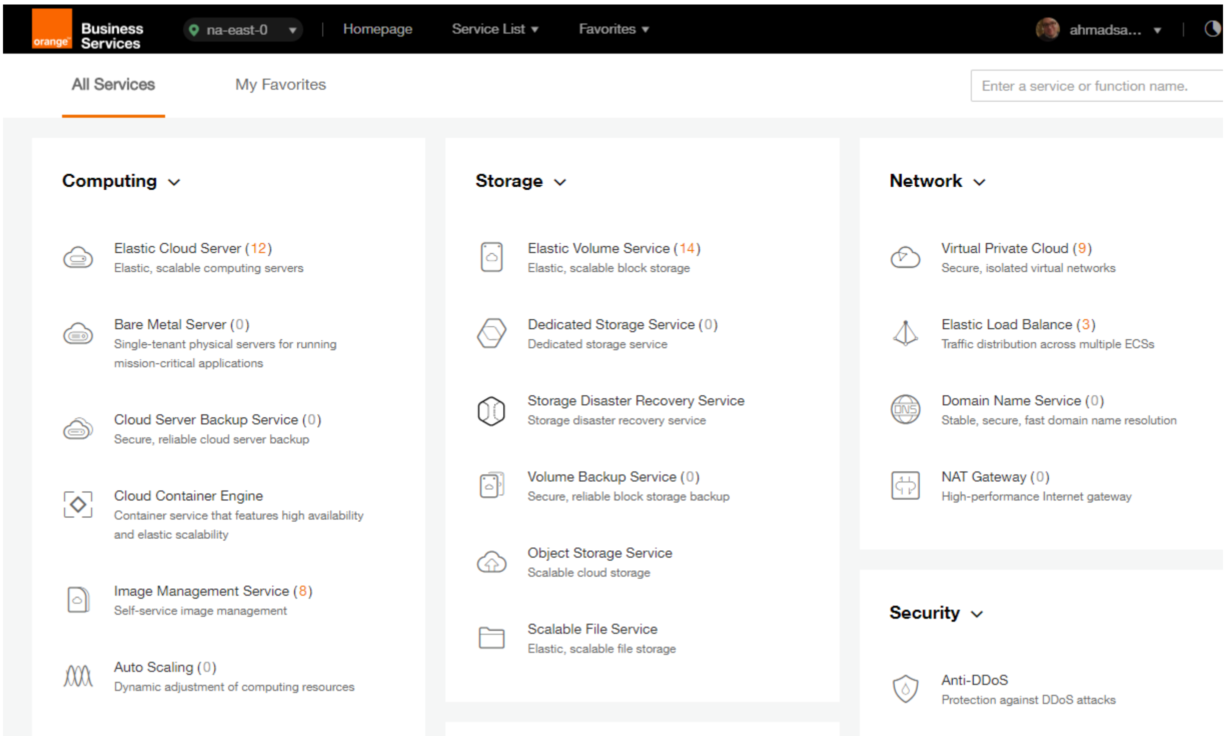

Business Services | na-east-0 | Homepage | Service List | Favorites | ahmadsa...

All Services | My Favorites | Enter a service or function name.

**Computing**

- Elastic Cloud Server (12) — Elastic, scalable computing servers
- Bare Metal Server (0) — Single-tenant physical servers for running mission-critical applications
- Cloud Server Backup Service (0) — Secure, reliable cloud server backup
- Cloud Container Engine — Container service that features high availability and elastic scalability
- Image Management Service (8) — Self-service image management
- Auto Scaling (0) — Dynamic adjustment of computing resources

**Storage**

- Elastic Volume Service (14) — Elastic, scalable block storage
- Dedicated Storage Service (0) — Dedicated storage service
- Storage Disaster Recovery Service — Storage disaster recovery service
- Volume Backup Service (0) — Secure, reliable block storage backup
- Object Storage Service — Scalable cloud storage
- Scalable File Service — Elastic, scalable file storage

**Network**

- Virtual Private Cloud (9) — Secure, isolated virtual networks
- Elastic Load Balance (3) — Traffic distribution across multiple ECSs
- Domain Name Service (0) — Stable, secure, fast domain name resolution
- NAT Gateway (0) — High-performance Internet gateway

**Security**

- Anti-DDoS — Protection against DDoS attacks

3. On the **Dashboard** page, click **Create VPC**.

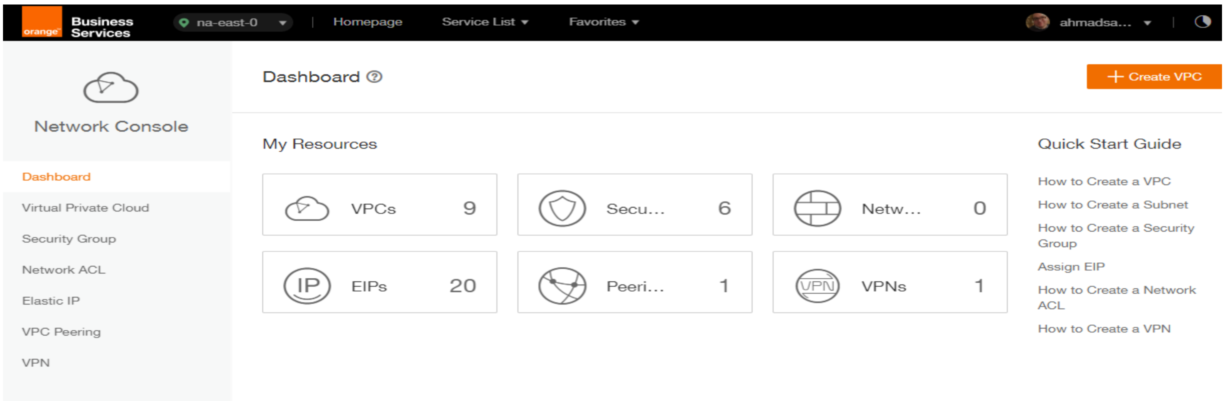

Business Services | na-east-0 | Homepage | Service List | Favorites | ahmadsa...

Network Console

- Dashboard
- Virtual Private Cloud
- Security Group
- Network ACL
- Elastic IP
- VPC Peering
- VPN

Dashboard | + Create VPC

My Resources

| | |
|---|---|
| VPCs | 9 |
| Secu... | 6 |
| Netw... | 0 |
| EIPs | 20 |
| Peeri... | 1 |
| VPNs | 1 |

Quick Start Guide

- How to Create a VPC
- How to Create a Subnet
- How to Create a Security Group
- Assign EIP
- How to Create a Network ACL
- How to Create a VPN

On the displayed **Apply for VPC** page, set the parameters as prompted.

Table 1 Parameter description

| Parameter | Description | Example Value |
|---|---|---|
| Name | Specifies the VPC name. | VPC-001 |
| VPC CIDR | Specifies the Classless Inter-Domain Routing (CIDR) block for the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset (for multiple subnets in the VPC). | 192.168.0.0/16 |

**Table 1** Parameter description

| Parameter | Description | Example Value |
|---|---|---|
| | The following CIDR blocks are supported:<br>10.0.0.0/8–24<br>172.16.0.0/12–24<br>192.168.0.0/16–24 | |
| Name | Specifies the subnet name. | Subnet-001 |
| CIDR | Specifies the CIDR block for the subnet. This value must be within the VPC CIDR range. | 192.168.0.0/24 |
| Gateway | Specifies the gateway address of the subnet. | 192.168.0.1 |

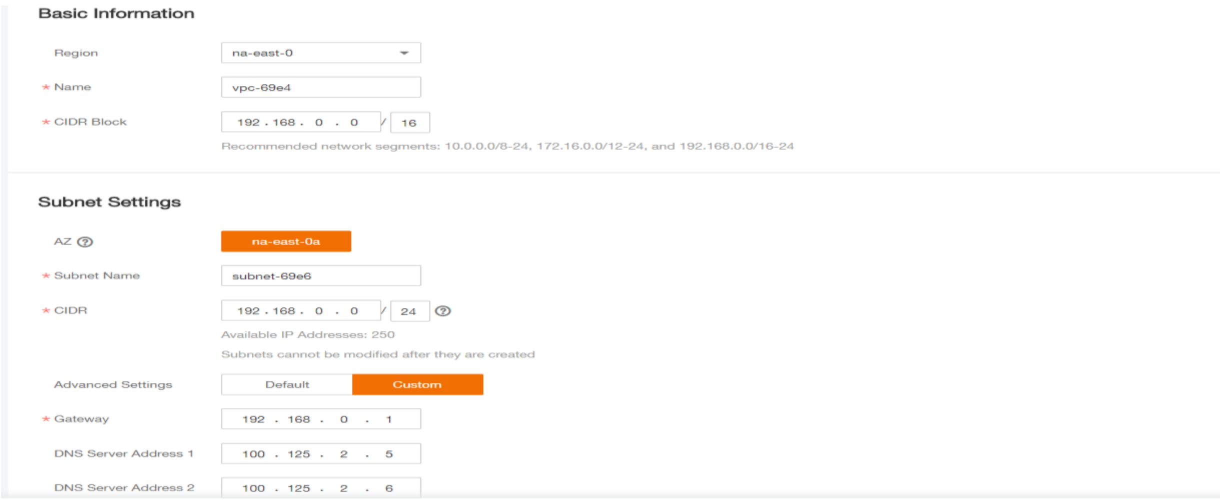

Basic Information

Region: na-east-0

* Name: vpc-69e4

* CIDR Block: 192 . 168 . 0 . 0 / 16

Recommended network segments: 10.0.0.0/8-24, 172.16.0.0/12-24, and 192.168.0.0/16-24

Subnet Settings

AZ: na-east-0a

* Subnet Name: subnet-69e6

* CIDR: 192 . 168 . 0 . 0 / 24

Available IP Addresses: 250

Subnets cannot be modified after they are created

Advanced Settings: Default [✓ Custom]

* Gateway: 192 . 168 . 0 . 1

DNS Server Address 1: 100 . 125 . 2 . 5

DNS Server Address 2: 100 . 125 . 2 . 6

4. The external DNS server address is used by default. If you need to change the DNS server address, click **Show Advanced Settings** and configure the DNS server addresses. You must ensure that the configured DNS server addresses are available.

5. Click **Create Now**.

   The created VPC will be shown in the VPC List

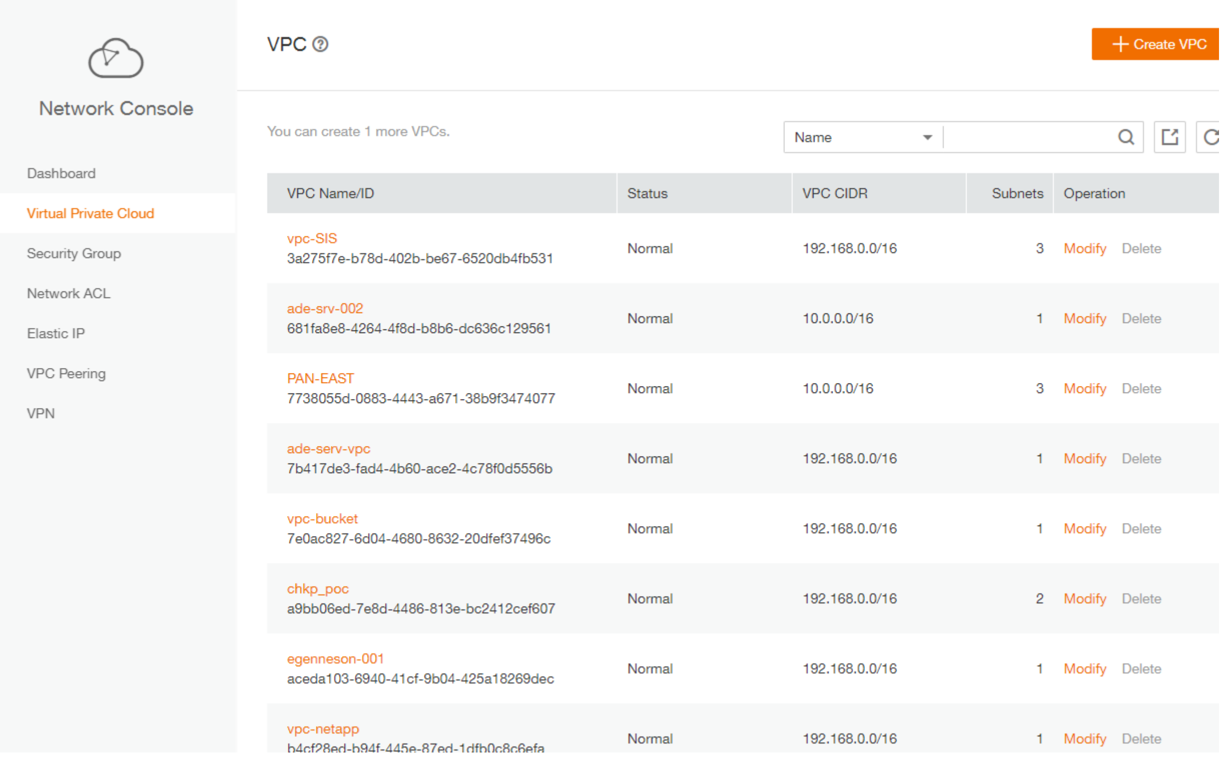

You can create 1 more VPCs.

| VPC Name/ID | Status | VPC CIDR | Subnets | Operation |
|---|---|---|---|---|
| vpc-SIS<br>3a275f7e-b78d-402b-be67-6520db4fb531 | Normal | 192.168.0.0/16 | 3 | Modify Delete |
| ade-srv-002<br>681fa8e8-4264-4f8d-b8b6-dc636c129561 | Normal | 10.0.0.0/16 | 1 | Modify Delete |
| PAN-EAST<br>7738055d-0883-4443-a671-38b9f3474077 | Normal | 10.0.0.0/16 | 3 | Modify Delete |
| ade-serv-vpc<br>7b417de3-fad4-4b60-ace2-4c78f0d5556b | Normal | 192.168.0.0/16 | 1 | Modify Delete |
| vpc-bucket<br>7e0ac827-6d04-4680-8632-20dfef37496c | Normal | 192.168.0.0/16 | 1 | Modify Delete |
| chkp_poc<br>a9bb06ed-7e8d-4486-813e-bc2412cef607 | Normal | 192.168.0.0/16 | 2 | Modify Delete |
| egenneson-001<br>aceda103-6940-41cf-9b04-425a18269dec | Normal | 192.168.0.0/16 | 1 | Modify Delete |
| vpc-netapp<br>b4cf28ed-b94f-445e-87ed-1dfb0c8c6efa | Normal | 192.168.0.0/16 | 1 | Modify Delete |

## 7.2 Install Palo Alto VM on the VPC

ECSs are more cost-effective than physical servers. Within minutes, you can obtain ECS resources from the public cloud. ECS resources are flexible and on-demand. This section describes how to create an ECS.

1. Log in to the management console.

2. Under **Computing**, click **Elastic Cloud Server**.

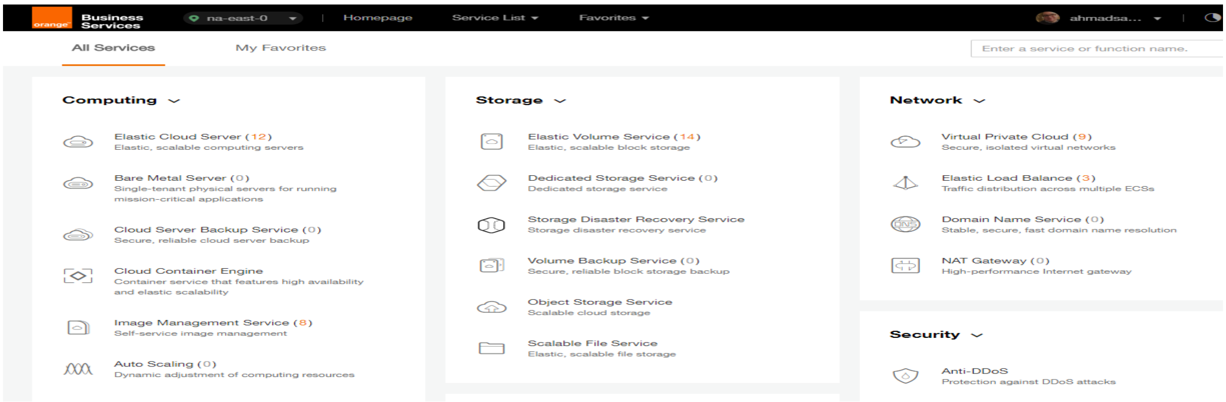

**Computing**

Elastic Cloud Server (12)
Elastic, scalable computing servers

Bare Metal Server (0)
Single-tenant physical servers for running mission-critical applications

Cloud Server Backup Service (0)
Secure, reliable cloud server backup

Cloud Container Engine
Container service that features high availability and elastic scalability

Image Management Service (8)
Self-service image management

Auto Scaling (0)
Dynamic adjustment of computing resources

**Storage**

Elastic Volume Service (14)
Elastic, scalable block storage

Dedicated Storage Service (0)
Dedicated storage service

Storage Disaster Recovery Service
Storage disaster recovery service

Volume Backup Service (0)
Secure, reliable block storage backup

Object Storage Service
Scalable cloud storage

Scalable File Service
Elastic, scalable file storage

**Network**

Virtual Private Cloud (9)
Secure, isolated virtual networks

Elastic Load Balance (3)
Traffic distribution across multiple ECSs

Domain Name Service (0)
Stable, secure, fast domain name resolution

NAT Gateway (0)
High-performance Internet gateway

**Security**

Anti-DDoS
Protection against DDoS attacks

3. Click **Create ECS**.

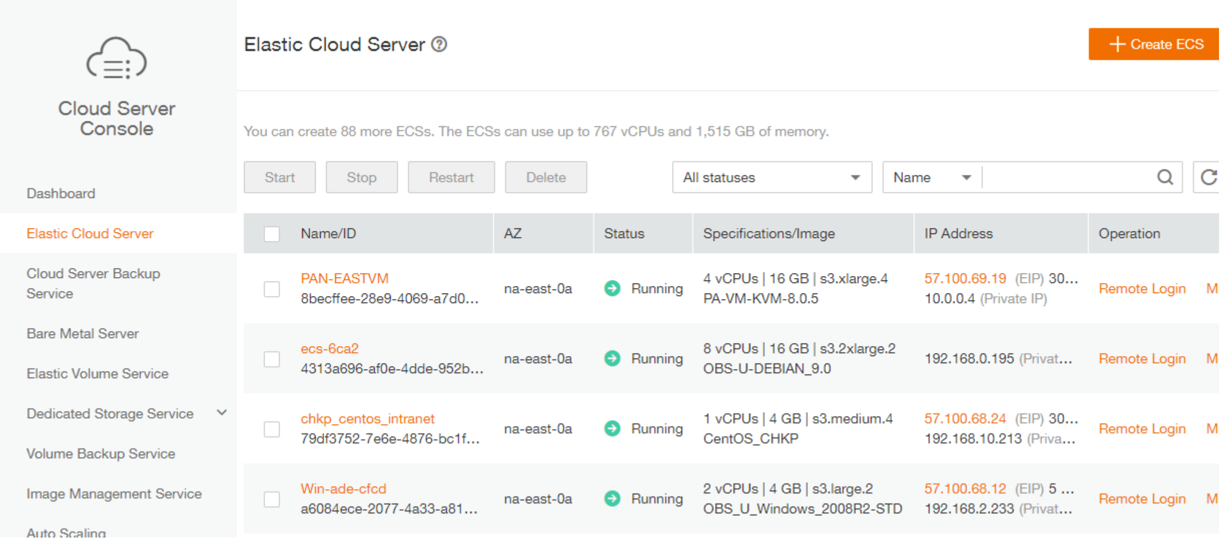

Elastic Cloud Server

You can create 88 more ECSs. The ECSs can use up to 767 vCPUs and 1,515 GB of memory.

| □ | Name/ID | AZ | Status | Specifications/Image | IP Address | Operation |
|---|---|---|---|---|---|---|
| □ | PAN-EASTVM<br>8becffee-28e9-4069-a7d0... | na-east-0a | Running | 4 vCPUs \| 16 GB \| s3.xlarge.4<br>PA-VM-KVM-8.0.5 | 57.100.69.19 (EIP) 30...<br>10.0.0.4 (Private IP) | Remote Login M |
| □ | ecs-6ca2<br>4313a696-af0e-4dde-952b... | na-east-0a | Running | 8 vCPUs \| 16 GB \| s3.2xlarge.2<br>OBS-U-DEBIAN_9.0 | 192.168.0.195 (Privat... | Remote Login M |
| □ | chkp_centos_intranet<br>79df3752-7e6e-4876-bc1f... | na-east-0a | Running | 1 vCPUs \| 4 GB \| s3.medium.4<br>CentOS_CHKP | 57.100.68.24 (EIP) 30...<br>192.168.10.213 (Priva... | Remote Login M |
| □ | Win-ade-cfcd<br>a6084ece-2077-4a33-a81... | na-east-0a | Running | 2 vCPUs \| 4 GB \| s3.large.2<br>OBS_U_Windows_2008R2-STD | 57.100.68.12 (EIP) 5 ...<br>192.168.2.233 (Privat... | Remote Login M |

The ECS creation page is displayed.

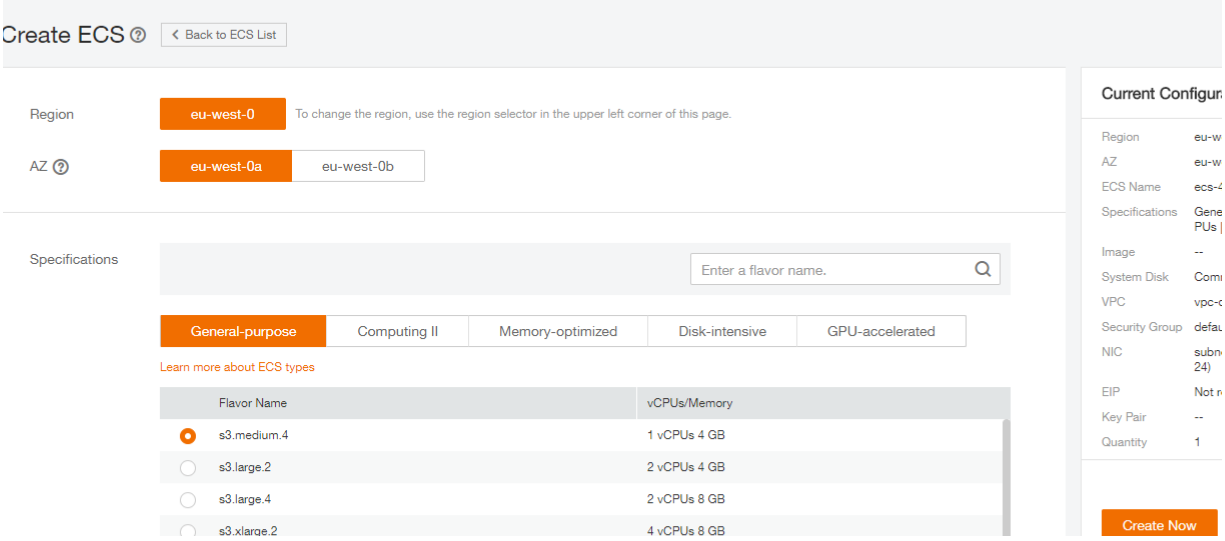

Region: eu-west-0 — To change the region, use the region selector in the upper left corner of this page.

AZ: eu-west-0a | eu-west-0b

Specifications: Enter a flavor name. — General-purpose

| | Flavor Name | vCPUs/Memory |
|---|---|---|
| ● | s3.medium.4 | 1 vCPUs 4 GB |
| ○ | s3.large.2 | 2 vCPUs 4 GB |
| ○ | s3.large.4 | 2 vCPUs 8 GB |
| ○ | s3.xlarge.2 | 4 vCPUs 8 GB |

Current Configur

| | |
|---|---|
| Region | eu-w |
| AZ | eu-w |
| ECS Name | ecs-4 |
| Specifications | Gene<br>PUs \| |
| Image | -- |
| System Disk | Comi |
| VPC | vpc-c |
| Security Group | defau |
| NIC | subn<br>24) |
| EIP | Not r |
| Key Pair | -- |
| Quantity | 1 |

4. Confirm the region.

If the region is incorrect, click 📍 in the upper left corner of the page for correction.

5. Select an AZ.

An AZ is a physical region where power and networks are physically isolated. AZs in the same region can communicate with each other over an intranet.

1. To enhance application availability, create ECSs in different AZs.
2. To shorten network latency, create ECSs in the same AZ.

6. Click + to open the **Select Specifications** page. On the page, select an ECS type.

7. Set **Local-Disk.**

This parameter is optional and is automatically displayed when you use a local disk.

A local disk specifies the local storage for the physical host where the ECS is deployed. Only hard disk drives (HDDs) are supported. If you select the disk-intensive ECS type, the system automatically attaches local disks to the ECS.

For example, if the Local Disk value is 3 x 1800 GB, three HDDs are attached to the ECS and the capacity of each HDD is 1800 GB.

8. Click **Image.**

Private Image

A private image is an image available only to the user who creates it. It contains an OS, preinstalled public applications, and the user's private applications. Using a private image to create ECSs removes the need to configure multiple ECSs repeatedly.

In our installation we previuosly uploaded a KVM image for PaloAlto VM . to check how to upload a private image to certain region please check the URL: [https://docs.prod-cloud-ocb.orange-business.com/en-us/usermanual/ims/en-us_topic_0030713190.html](https://docs.prod-cloud-ocb.orange-business.com/en-us/usermanual/ims/en-us_topic_0030713190.html)

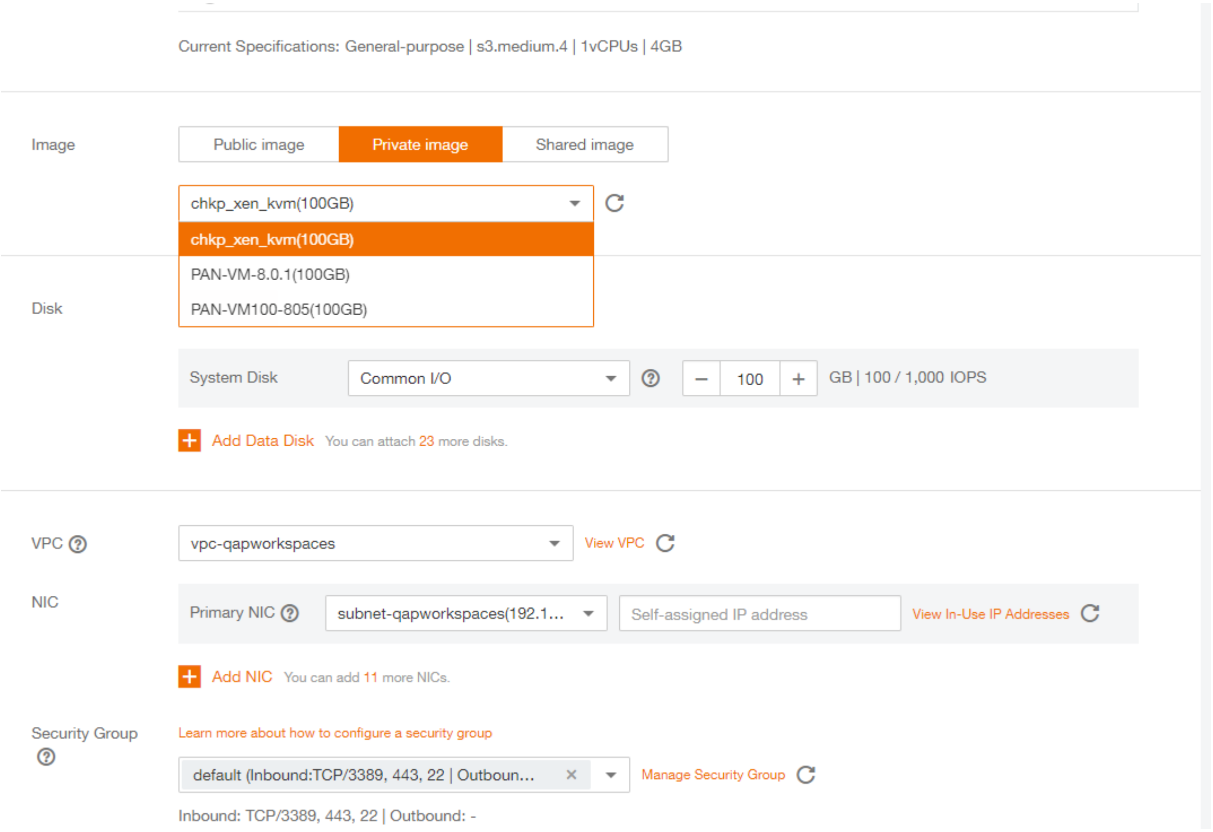

Current Specifications: General-purpose | s3.medium.4 | 1vCPUs | 4GB

Image: Public image | Private image | Shared image

chkp_xen_kvm(100GB)

chkp_xen_kvm(100GB)
PAN-VM-8.0.1(100GB)
PAN-VM100-805(100GB)

Disk: System Disk Common I/O 100 GB | 100 / 1,000 IOPS

Add Data Disk You can attach 23 more disks.

VPC: vpc-qapworkspaces View VPC

NIC: Primary NIC subnet-qapworkspaces(192.1... Self-assigned IP address View In-Use IP Addresses

Add NIC You can add 11 more NICs.

Security Group: Learn more about how to configure a security group

default (Inbound:TCP/3389, 443, 22 | Outboun... Manage Security Group

Inbound: TCP/3389, 443, 22 | Outbound: -

9. Set **Disk.**

A disk can be a system disk or a data disk. You can create multiple data disks for an ECS and customize their disk sizes.

10. Set network parameters, including VPC, Security Group, and NIC.

When you use VPC for the first time, the system automatically creates a VPC for you, including the security group and NIC.

**Table 2** Parameter descriptions

| Parameter | Description |
|---|---|
| VPC | Provides a network, including subnet and security group, for an ECS.<br>You can select an existing VPC, or click **View VPC** and create a desired one.<br>For more information about VPC, see *Virtual Private Cloud User Guide*.<br>**NOTE:**<br>DHCP must be enabled in the VPC to which the ECS belongs. |
| Security Group | Controls instance access within or between security groups by defining access rules. This enhances instance security.<br>When creating an ECS, you can select multiple (recommended not more than five) security groups. In such a case, the access rules of all the selected security groups apply on the ECS.<br>**NOTE:**<br>Before initializing an ECS, ensure that the security group rule in the outbound direction meets the following requirements:<br>• **Protocol**: **TCP**<br>• **Port Range**: **80**<br>• **Remote End**: **169.254.0.0/16**<br>If you use the default security group rule in the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rule in the outbound direction is as follows:<br>• **Protocol**: **ANY**<br>• **Port Range**: **ANY**<br>• **Remote End**: **0.0.0.0/16** |
| NIC | Consists of a primary NIC and one or more extension NICs.<br>**MTU Settings**: optional<br>If your ECS is of M2, large-memory, H1, or D1 type, you can click **MTU Settings** to configure the maximum transmission unit (MTU) for a to-be-added extension NIC for improving network performance.<br>An MTU can only be a number, ranging from 1280 to 8888.<br>** In our scenario: We created only two NIC cards one for the Management and the Other is for the Untrust Interfaces.<br>The other two NIC cards will be created using API request on the Business and Web VPC's then will be assigned to the Palo Alto |

**Table 2** Parameter descriptions

| Parameter | Description |
|---|---|
| | VM ** |
| EIP | A static public IP address bound to an ECS in a VPC. Using the EIP, the ECS provides services externally.<br>The following options are provided:<br>• **Do not use**<br>Without an EIP, the ECS cannot access the Internet and is used only in the private network or cluster.<br>• **Automatically assign**<br>The system automatically assigns an EIP for the ECS. The EIP provides exclusive bandwidth that is configurable.<br>• **Specify**<br>An existing EIP is assigned for the ECS. When using an existing EIP, you cannot create ECSs in batches.<br>** In our scenario: We assigned 2 EIP's one for the management NIC and the other for the Un trust NIC. |

11. Set ECS Name.

    If you want to create multiple ECSs at a time, the system automatically sequences these ECSs.

12. Configure the number of ECSs to be created.

    After the configuration, click Price Calculator to view the ECS configuration fee.

13. Click Create Now.

14. On the ECS specification confirmation page, confirm the ECS specifications and click Submit.

    After the ECS is created, you can view information about it on the Elastic Cloud Server page.

15. After creating the Palo Alto VM you can access it through Https using the EIP of the Management NIC. (username: admin / Password: admin)

# 8 Palo Alto VM-Series High Availability configuration

## 8.1 Adding Virtual IP’s on the Un-trust and Trust Subnets

1. Login to the management console
2. From Network - Choose virtual private cloud

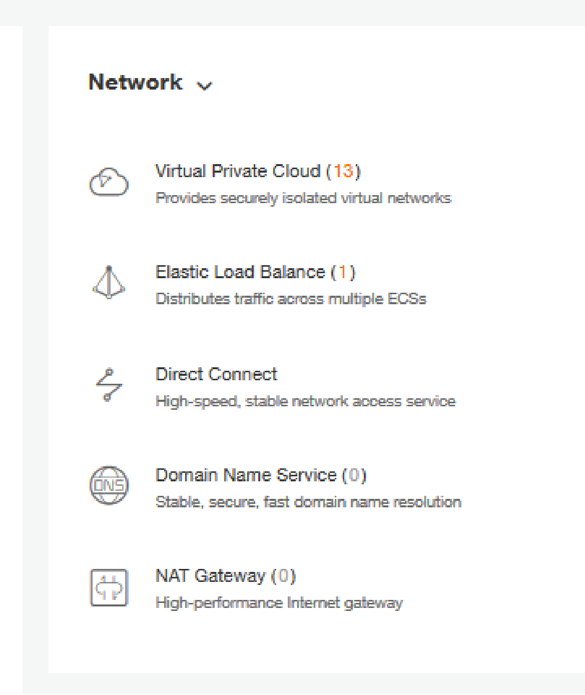

Network

Virtual Private Cloud (13)
Provides securely isolated virtual networks

Elastic Load Balance (1)
Distributes traffic across multiple ECSs

Direct Connect
High-speed, stable network access service

Domain Name Service (0)
Stable, secure, fast domain name resolution

NAT Gateway (0)
High-performance Internet gateway

3. Choose the VPC hosting the Palo Alto VM’s and click on Subnets

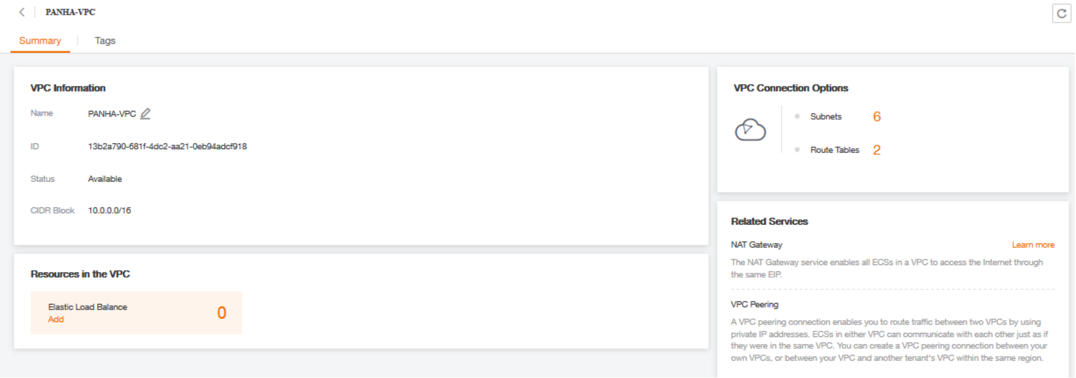

PANHA-VPC

Summary | Tags

VPC Information

| | |
|---|---|
| Name | PANHA-VPC |
| ID | 13b2a790-681f-4dc2-aa21-0eb94adcf918 |
| Status | Available |
| CIDR Block | 10.0.0.0/16 |

Resources in the VPC

Elastic Load Balance 0
Add

VPC Connection Options

| | |
|---|---|
| Subnets | 6 |
| Route Tables | 2 |

Related Services

NAT Gateway Learn more

The NAT Gateway service enables all ECSs in a VPC to access the Internet through the same EIP.

VPC Peering

A VPC peering connection enables you to route traffic between two VPCs by using private IP addresses. ECSs in either VPC can communicate with each other just as if they were in the same VPC. You can create a VPC peering connection between your own VPCs, or between your VPC and another tenant's VPC within the same region.

4. Choose the Un-trust Subnet

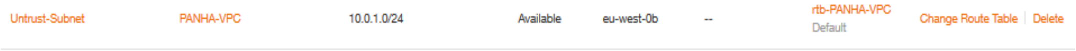

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| Untrust-Subnet | PANHA-VPC | 10.0.1.0/24 | Available | eu-west-0b | -- | rtb-PANHA-VPC<br>Default | Change Route Table \| Delete |

5. Click on IP Addresses Tab

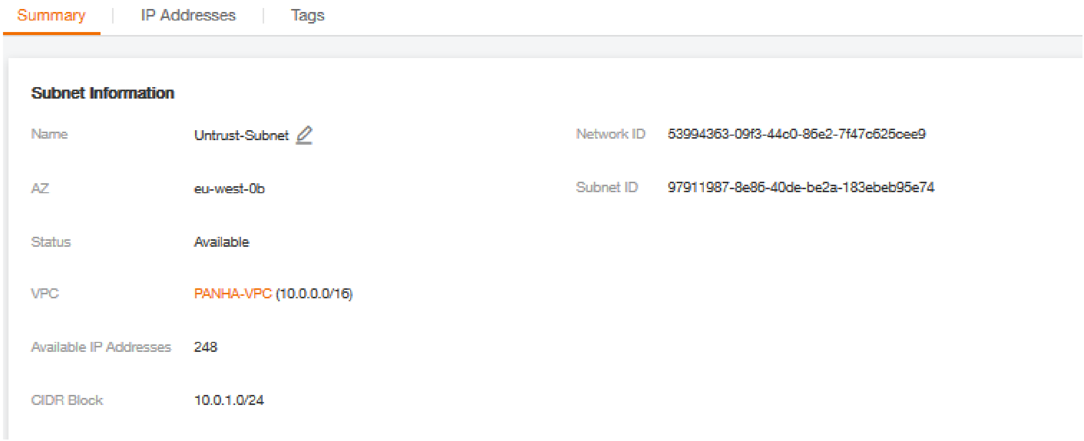

**Subnet Information**

| | | | |
|---|---|---|---|
| Name | Untrust-Subnet | Network ID | 53994363-09f3-44c0-86e2-7f47c625cee9 |
| AZ | eu-west-0b | Subnet ID | 97911987-8e86-40de-be2a-183ebeb95e74 |
| Status | Available | | |
| VPC | PANHA-VPC (10.0.0.0/16) | | |
| Available IP Addresses | 248 | | |
| CIDR Block | 10.0.1.0/24 | | |

6. Assign Virtual IP Address

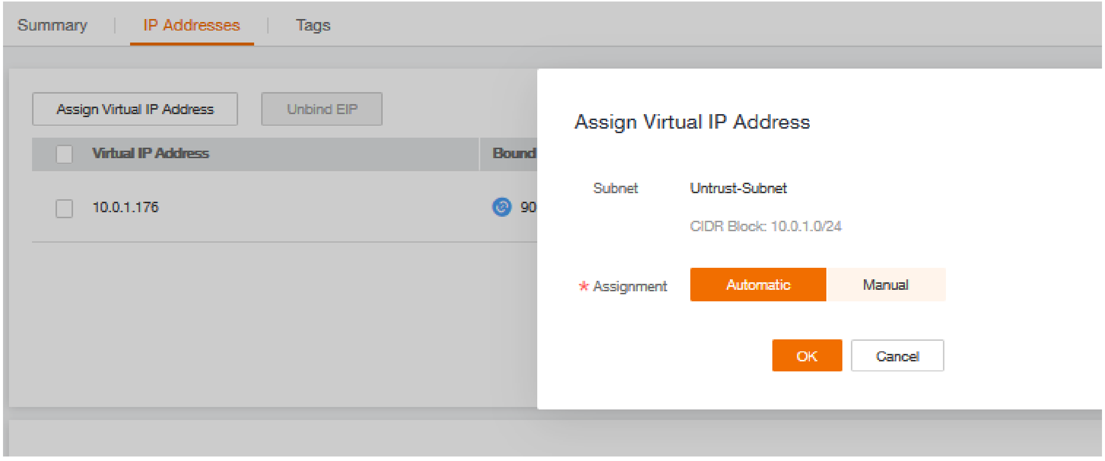

| Virtual IP Address | Bound |
|---|---|
| 10.0.1.176 | 90 |

Assign Virtual IP Address

Subnet Untrust-Subnet

CIDR Block: 10.0.1.0/24

* Assignment Automatic

7. Assign EIP to the Virtual IP created.

8. Bind the VIP the un-trust ports of the primary and backup VM-Series firewall.

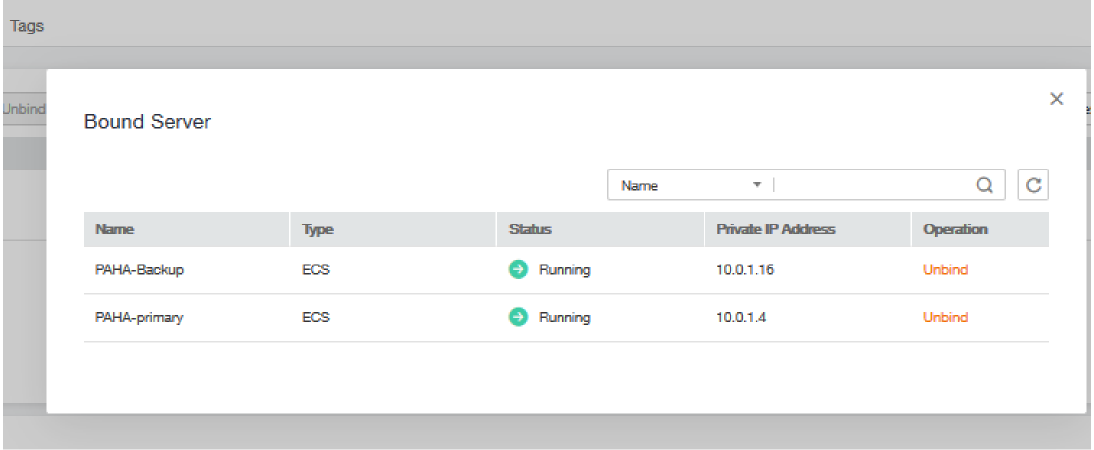

Bound Server

| Name | Type | Status | Private IP Address | Operation |
|---|---|---|---|---|
| PAHA-Backup | ECS | Running | 10.0.1.16 | Unbind |
| PAHA-primary | ECS | Running | 10.0.1.4 | Unbind |

9. Perform the Same steps with the **Trust subnet** but don't assign an EIp to the virtual IP in the trust Subnet.

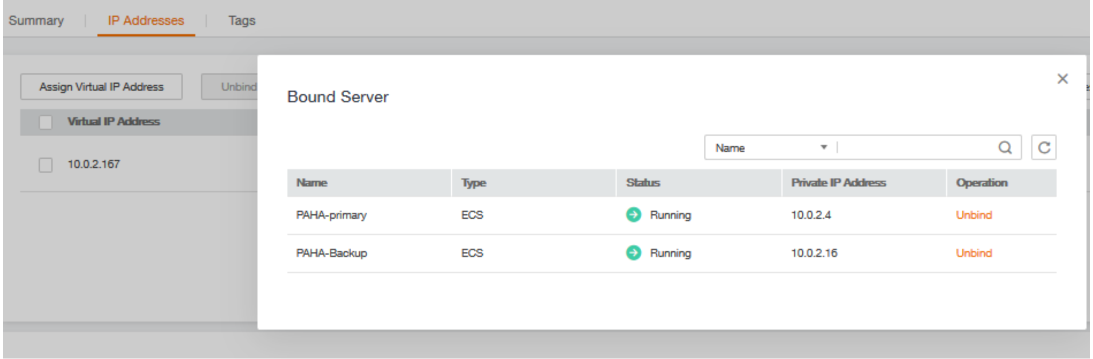

| Virtual IP Address |
|---|
| 10.0.2.167 |

Bound Server

| Name | Type | Status | Private IP Address | Operation |
|---|---|---|---|---|
| PAHA-primary | ECS | Running | 10.0.2.4 | Unbind |
| PAHA-Backup | ECS | Running | 10.0.2.16 | Unbind |

9.Uncheck the Source/Destination Check for all interfaces of each firewall except the one for the admin NIC interface

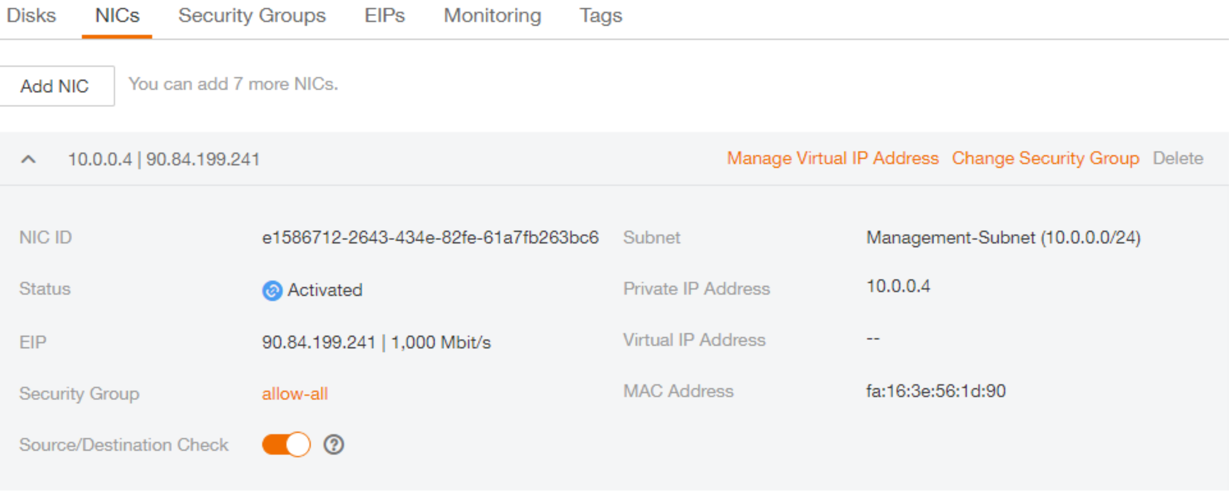

You can add 7 more NICs.

10.0.0.4 | 90.84.199.241

|  |  |  |  |
|---|---|---|---|
| NIC ID | e1586712-2643-434e-82fe-61a7fb263bc6 | Subnet | Management-Subnet (10.0.0.0/24) |
| Status | Activated | Private IP Address | 10.0.0.4 |
| EIP | 90.84.199.241 \| 1,000 Mbit/s | Virtual IP Address | -- |
| Security Group | allow-all | MAC Address | fa:16:3e:56:1d:90 |
| Source/Destination Check | ● |  |  |

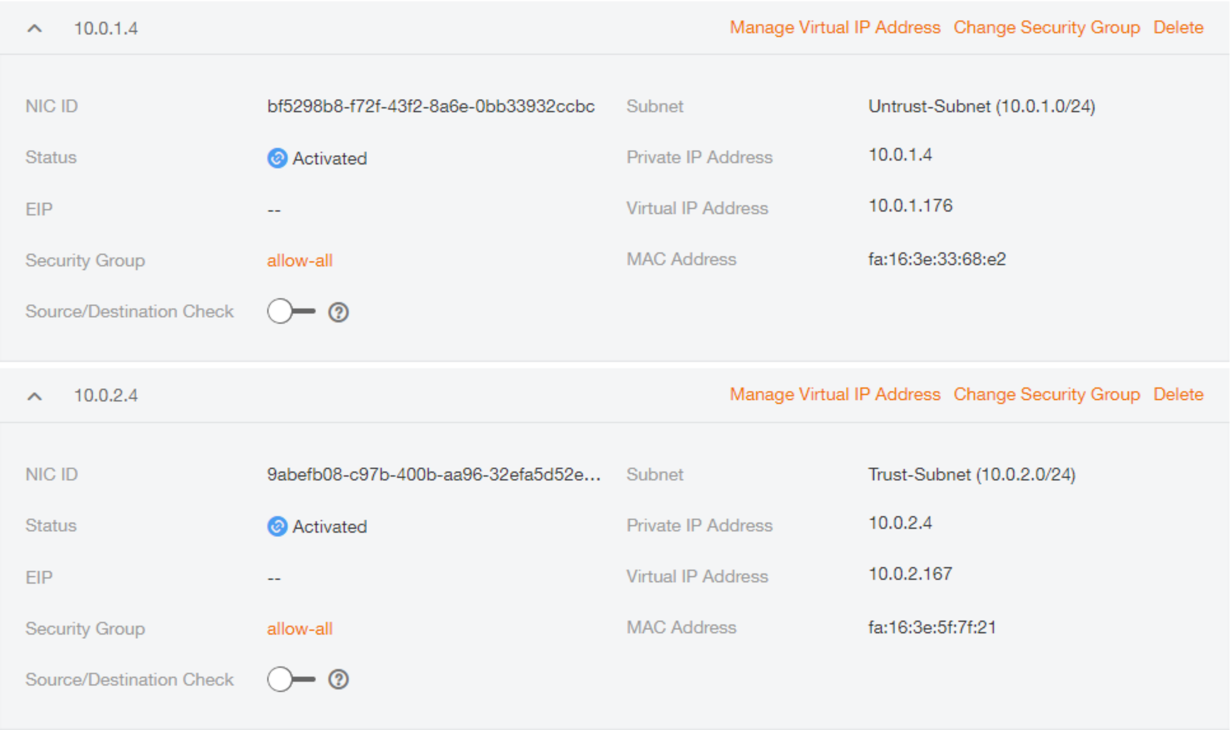

^ 10.0.1.4 Manage Virtual IP Address Change Security Group Delete

| | | | |
|---|---|---|---|
| NIC ID | bf5298b8-f72f-43f2-8a6e-0bb33932ccbc | Subnet | Untrust-Subnet (10.0.1.0/24) |
| Status | Activated | Private IP Address | 10.0.1.4 |
| EIP | -- | Virtual IP Address | 10.0.1.176 |
| Security Group | allow-all | MAC Address | fa:16:3e:33:68:e2 |
| Source/Destination Check | ○ ? | | |

^ 10.0.2.4 Manage Virtual IP Address Change Security Group Delete

| | | | |
|---|---|---|---|
| NIC ID | 9abefb08-c97b-400b-aa96-32efa5d52e... | Subnet | Trust-Subnet (10.0.2.0/24) |
| Status | Activated | Private IP Address | 10.0.2.4 |
| EIP | -- | Virtual IP Address | 10.0.2.167 |
| Security Group | allow-all | MAC Address | fa:16:3e:5f:7f:21 |
| Source/Destination Check | ○ ? | | |

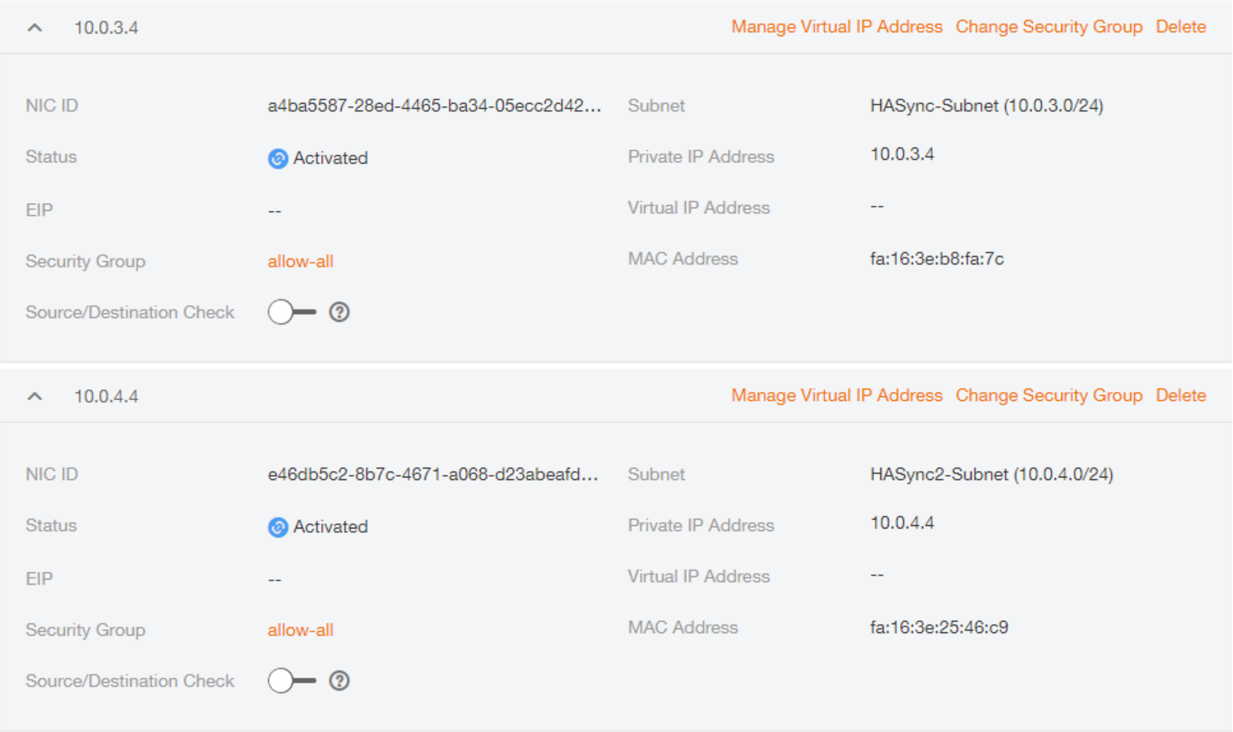

^ 10.0.3.4 Manage Virtual IP Address Change Security Group Delete

| | | | |
|---|---|---|---|
| NIC ID | a4ba5587-28ed-4465-ba34-05ecc2d42... | Subnet | HASync-Subnet (10.0.3.0/24) |
| Status | Activated | Private IP Address | 10.0.3.4 |
| EIP | -- | Virtual IP Address | -- |
| Security Group | allow-all | MAC Address | fa:16:3e:b8:fa:7c |
| Source/Destination Check | ○ ? | | |

^ 10.0.4.4 Manage Virtual IP Address Change Security Group Delete

| | | | |
|---|---|---|---|
| NIC ID | e46db5c2-8b7c-4671-a068-d23abeafd... | Subnet | HASync2-Subnet (10.0.4.0/24) |
| Status | Activated | Private IP Address | 10.0.4.4 |
| EIP | -- | Virtual IP Address | -- |
| Security Group | allow-all | MAC Address | fa:16:3e:25:46:c9 |
| Source/Destination Check | ○ ? | | |

10. Change the security Group for all interfaces to allow-all traffic , as security will be done through the firewall not the platform .

## 8.2 Initial configuration for the primary Palo Alto VM

1. Login the primary Palo Alto VM using the management port EIP through https

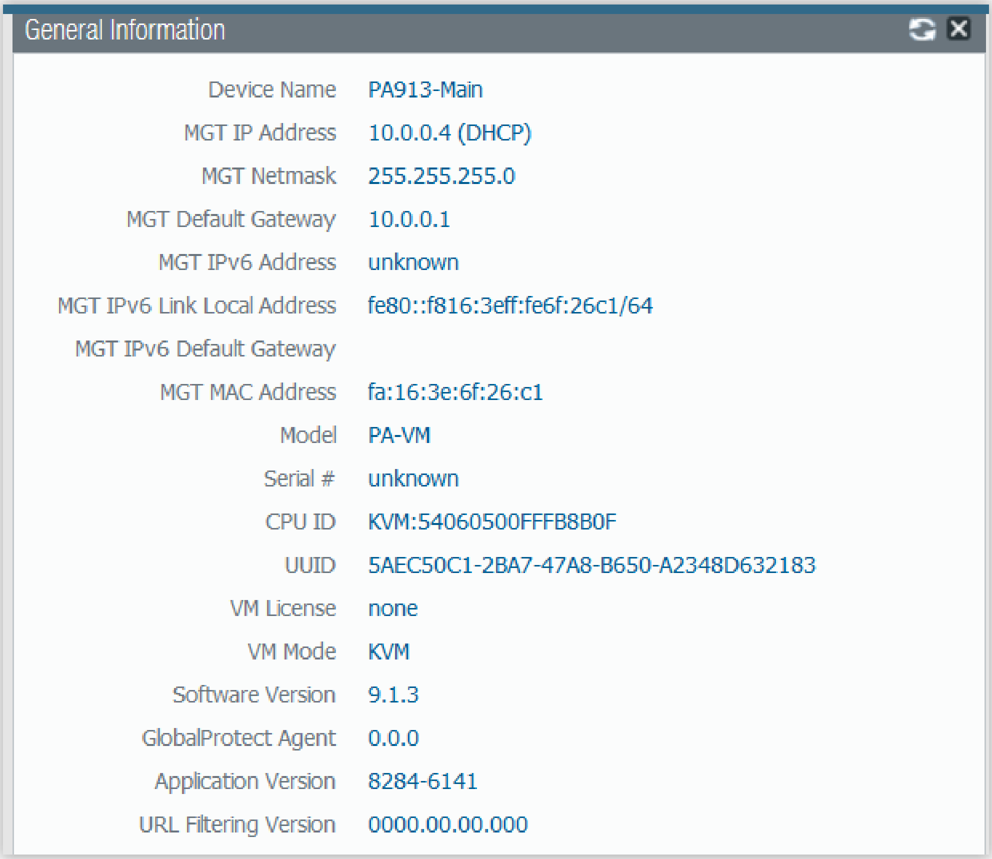

General Information

| | |
|---|---|
| Device Name | PA913-Main |
| MGT IP Address | 10.0.0.4 (DHCP) |
| MGT Netmask | 255.255.255.0 |
| MGT Default Gateway | 10.0.0.1 |
| MGT IPv6 Address | unknown |
| MGT IPv6 Link Local Address | fe80::f816:3eff:fe6f:26c1/64 |
| MGT IPv6 Default Gateway | |
| MGT MAC Address | fa:16:3e:6f:26:c1 |
| Model | PA-VM |
| Serial # | unknown |
| CPU ID | KVM:54060500FFFB8B0F |
| UUID | 5AEC50C1-2BA7-47A8-B650-A2348D632183 |
| VM License | none |
| VM Mode | KVM |
| Software Version | 9.1.3 |
| GlobalProtect Agent | 0.0.0 |
| Application Version | 8284-6141 |
| URL Filtering Version | 0000.00.00.000 |

2. Got to Network > Interfaces and configure the untrust and trust ports by adding the Virtual IP's you assigned on the Untrust and trust subnets.

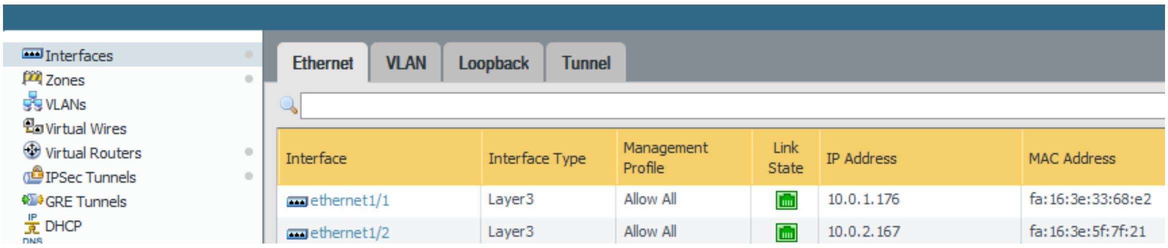

| Interface | Interface Type | Management Profile | Link State | IP Address | MAC Address |
|---|---|---|---|---|---|
| ethernet1/1 | Layer3 | Allow All | 🟩 | 10.0.1.176 | fa:16:3e:33:68:e2 |
| ethernet1/2 | Layer3 | Allow All | 🟩 | 10.0.2.167 | fa:16:3e:5f:7f:21 |

3. Commit

4. Go to Policies > Security > Add two policies between un-trust to trust and vice versa

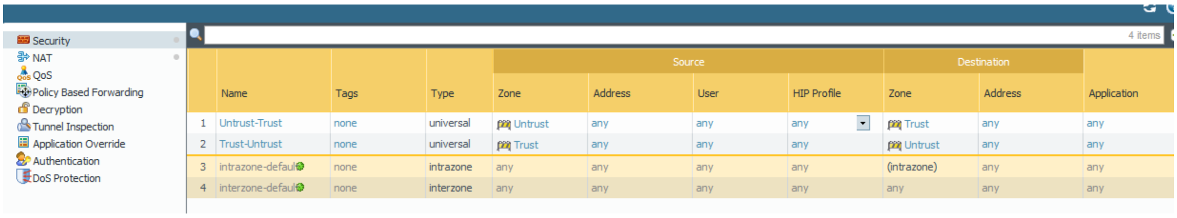

| | | | | Source | | | | Destination | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | Name | Tags | Type | Zone | Address | User | HIP Profile | Zone | Address | Application |
| 1 | Untrust-Trust | none | universal | Untrust | any | any | any | Trust | any | any |
| 2 | Trust-Untrust | none | universal | Trust | any | any | any | Untrust | any | any |
| 3 | intrazone-defaul | none | intrazone | any | any | any | any | (intrazone) | any | any |
| 4 | interzone-defaul | none | interzone | any | any | any | any | any | any | any |

5. Go to Policies > NAT > Add two NAT rules as shown below

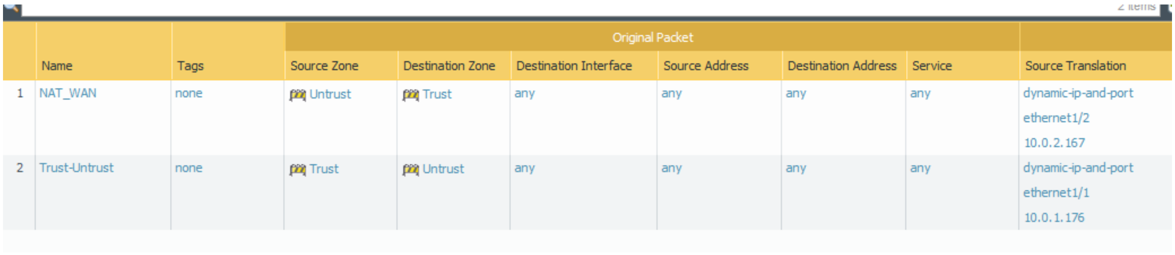

| | Name | Tags | Original Packet | | | | | | Source Translation |
|---|---|---|---|---|---|---|---|---|---|
| | | | Source Zone | Destination Zone | Destination Interface | Source Address | Destination Address | Service | |
| 1 | NAT_WAN | none | Untrust | Trust | any | any | any | any | dynamic-ip-and-port<br>ethernet1/2<br>10.0.2.167 |
| 2 | Trust-Untrust | none | Trust | Untrust | any | any | any | any | dynamic-ip-and-port<br>ethernet1/1<br>10.0.1.176 |

6. Add routes to enable traffic from untrust and protected zone and vice versa

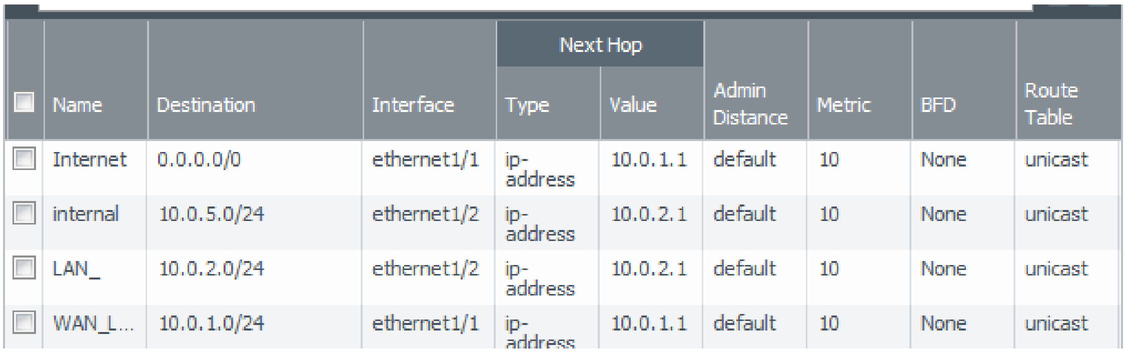

| □ | Name | Destination | Interface | Next Hop | | Admin Distance | Metric | BFD | Route Table |
|---|---|---|---|---|---|---|---|---|---|
| | | | | Type | Value | | | | |
| □ | Internet | 0.0.0.0/0 | ethernet1/1 | ip-address | 10.0.1.1 | default | 10 | None | unicast |
| □ | internal | 10.0.5.0/24 | ethernet1/2 | ip-address | 10.0.2.1 | default | 10 | None | unicast |
| □ | LAN_ | 10.0.2.0/24 | ethernet1/2 | ip-address | 10.0.2.1 | default | 10 | None | unicast |
| □ | WAN_L... | 10.0.1.0/24 | ethernet1/1 | ip-address | 10.0.1.1 | default | 10 | None | unicast |

7. Configure the High Availability Ports . In our scenario we have ports 3 and 4 .

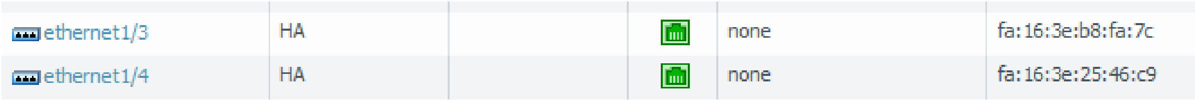

| | | | | | |
|---|---|---|---|---|---|
| ethernet1/3 | HA | | 🟩 | none | fa:16:3e:b8:fa:7c |
| ethernet1/4 | HA | | 🟩 | none | fa:16:3e:25:46:c9 |

8. Perform the same configuration to the backup firewall and make sure that the two firewalls have the same software version and identical to each other.

9. Now we have the 2 firewalls identically configured . They are ready to start the high availability configuration and synchronize with each other.

## 8.3 Firewall Software Update to release 9.1.4

From Device Tab > Choose Software

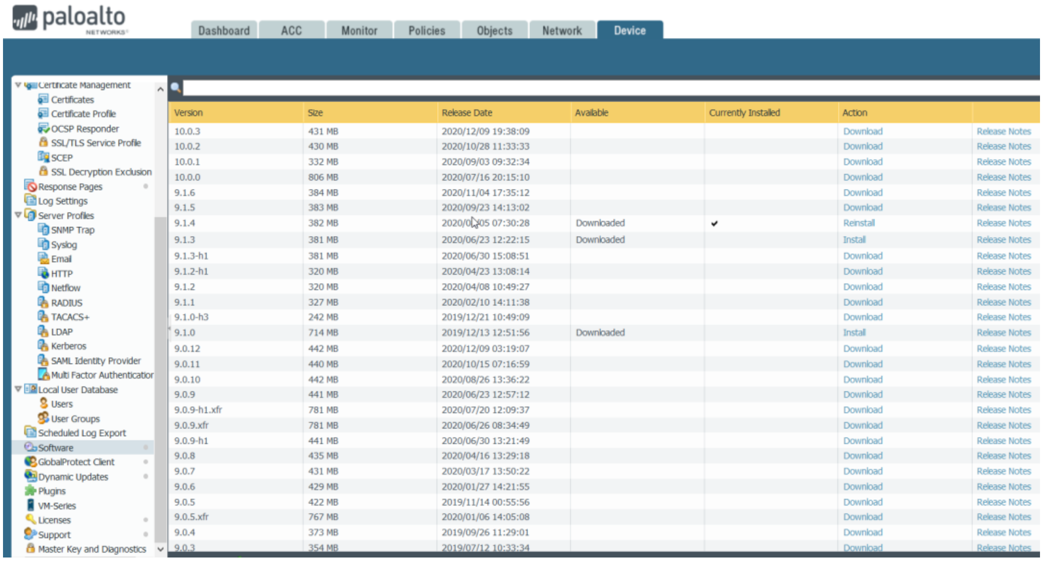

| Version | Size | Release Date | Available | Currently Installed | Action | |
|---|---|---|---|---|---|---|
| 10.0.3 | 431 MB | 2020/12/09 19:38:09 | | | Download | Release Notes |
| 10.0.2 | 430 MB | 2020/10/28 11:33:33 | | | Download | Release Notes |
| 10.0.1 | 332 MB | 2020/09/03 09:32:34 | | | Download | Release Notes |
| 10.0.0 | 806 MB | 2020/07/16 20:15:10 | | | Download | Release Notes |
| 9.1.6 | 384 MB | 2020/11/04 17:35:12 | | | Download | Release Notes |
| 9.1.5 | 383 MB | 2020/09/23 14:13:02 | | | Download | Release Notes |
| 9.1.4 | 382 MB | 2020/08/05 07:30:28 | Downloaded | ✓ | Reinstall | Release Notes |
| 9.1.3 | 381 MB | 2020/06/23 12:22:15 | Downloaded | | Install | Release Notes |
| 9.1.3-h1 | 381 MB | 2020/06/30 15:08:51 | | | Download | Release Notes |
| 9.1.2-h1 | 320 MB | 2020/04/23 13:08:14 | | | Download | Release Notes |
| 9.1.2 | 320 MB | 2020/04/08 10:49:27 | | | Download | Release Notes |
| 9.1.1 | 327 MB | 2020/02/10 14:11:38 | | | Download | Release Notes |
| 9.1.0-h3 | 242 MB | 2019/12/21 10:49:09 | | | Download | Release Notes |
| 9.1.0 | 714 MB | 2019/12/13 12:51:56 | Downloaded | | Install | Release Notes |
| 9.0.12 | 442 MB | 2020/12/09 03:19:07 | | | Download | Release Notes |
| 9.0.11 | 440 MB | 2020/10/15 07:16:59 | | | Download | Release Notes |
| 9.0.10 | 442 MB | 2020/08/26 13:36:22 | | | Download | Release Notes |
| 9.0.9 | 441 MB | 2020/06/23 12:57:12 | | | Download | Release Notes |
| 9.0.9-h1.xfr | 781 MB | 2020/07/20 12:09:37 | | | Download | Release Notes |
| 9.0.9.xfr | 781 MB | 2020/06/26 08:34:49 | | | Download | Release Notes |
| 9.0.9-h1 | 441 MB | 2020/06/30 13:21:49 | | | Download | Release Notes |
| 9.0.8 | 435 MB | 2020/04/16 13:29:18 | | | Download | Release Notes |
| 9.0.7 | 431 MB | 2020/03/17 13:50:22 | | | Download | Release Notes |
| 9.0.6 | 429 MB | 2020/01/27 14:21:55 | | | Download | Release Notes |
| 9.0.5 | 422 MB | 2019/11/14 00:55:56 | | | Download | Release Notes |
| 9.0.5.xfr | 767 MB | 2020/01/06 14:05:08 | | | Download | Release Notes |
| 9.0.4 | 373 MB | 2019/09/26 11:29:01 | | | Download | Release Notes |
| 9.0.3 | 354 MB | 2019/07/12 10:33:34 | | | Download | Release Notes |

Choose release 9.1.4 update > then click Download

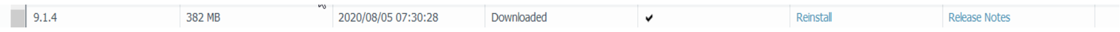

| | | | | | | |
|---|---|---|---|---|---|---|
| 9.1.4 | 382 MB | 2020/08/05 07:30:28 | Downloaded | ✓ | Reinstall | Release Notes |

The update will start to download after the download is finished . Click on install.

This will update the firwall software from 9.1.3 to 9.1.4

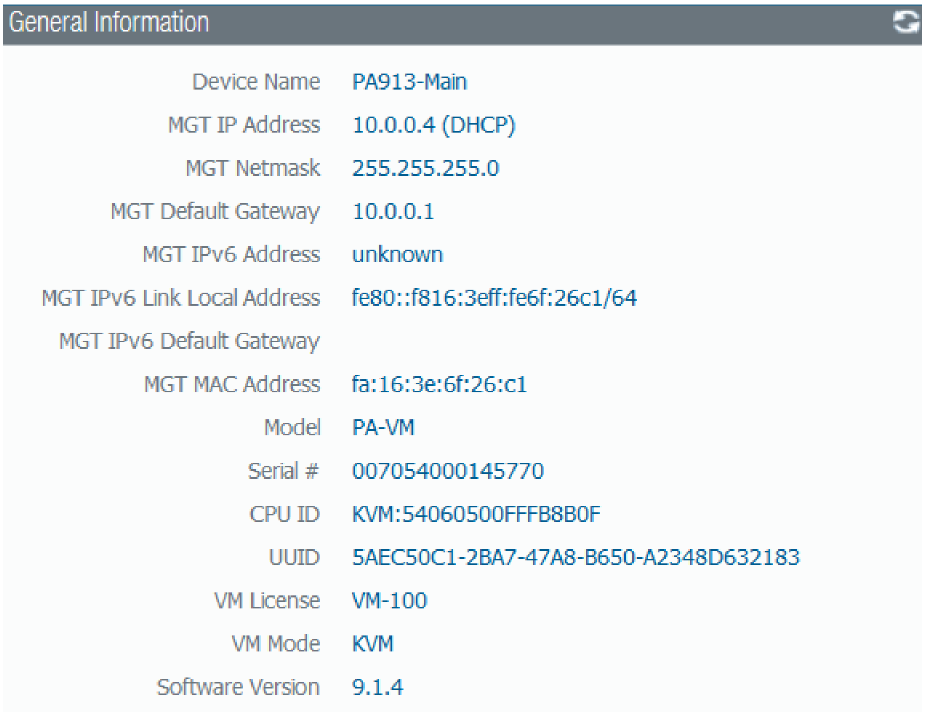

General Information

| | |
|---|---|
| Device Name | PA913-Main |
| MGT IP Address | 10.0.0.4 (DHCP) |
| MGT Netmask | 255.255.255.0 |
| MGT Default Gateway | 10.0.0.1 |
| MGT IPv6 Address | unknown |
| MGT IPv6 Link Local Address | fe80::f816:3eff:fe6f:26c1/64 |
| MGT IPv6 Default Gateway | |
| MGT MAC Address | fa:16:3e:6f:26:c1 |
| Model | PA-VM |
| Serial # | 007054000145770 |
| CPU ID | KVM:54060500FFFB8B0F |
| UUID | 5AEC50C1-2BA7-47A8-B650-A2348D632183 |
| VM License | VM-100 |
| VM Mode | KVM |
| Software Version | 9.1.4 |

## 8.4 High Availability (Active-Passive) Configuration

1. Choose Device > High Availability

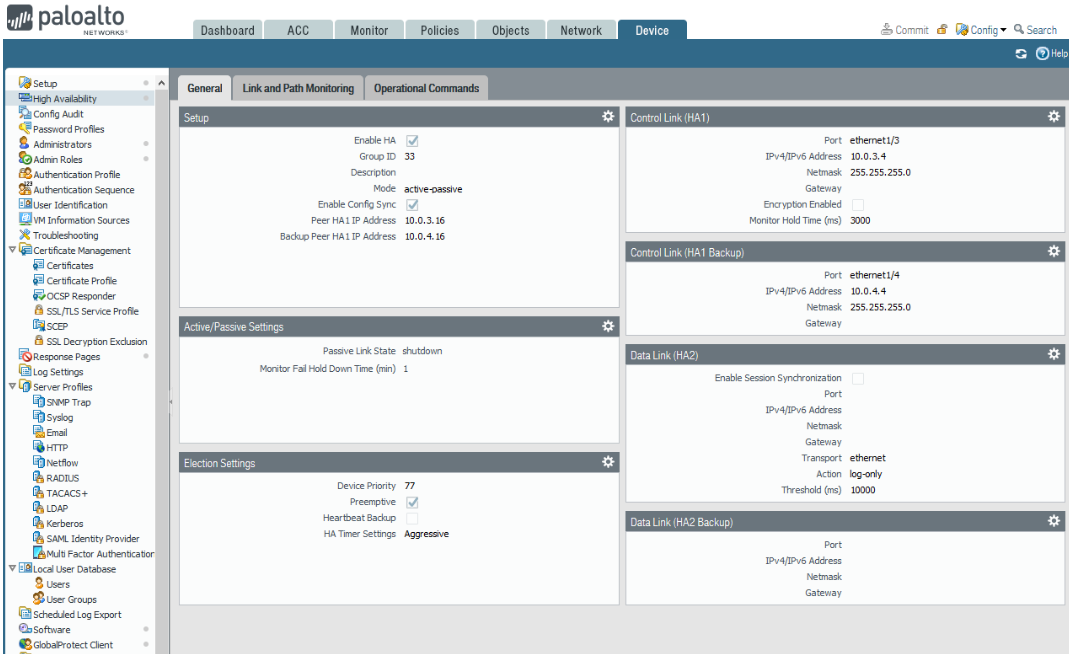

Setup

Enable HA ☑
Group ID 33
Description
Mode active-passive
Enable Config Sync ☑
Peer HA1 IP Address 10.0.3.16
Backup Peer HA1 IP Address 10.0.4.16

Active/Passive Settings

Passive Link State shutdown
Monitor Fail Hold Down Time (min) 1

Election Settings

Device Priority 77
Preemptive ☑
Heartbeat Backup □
HA Timer Settings Aggressive

Control Link (HA1)

Port ethernet1/3
IPv4/IPv6 Address 10.0.3.4
Netmask 255.255.255.0
Gateway
Encryption Enabled □
Monitor Hold Time (ms) 3000

Control Link (HA1 Backup)

Port ethernet1/4
IPv4/IPv6 Address 10.0.4.4
Netmask 255.255.255.0
Gateway

Data Link (HA2)

Enable Session Synchronization □
Port
IPv4/IPv6 Address
Netmask
Gateway
Transport ethernet
Action log-only
Threshold (ms) 10000

Data Link (HA2 Backup)

Port
IPv4/IPv6 Address
Netmask
Gateway

2. Form General Tab Choose Setup > settings

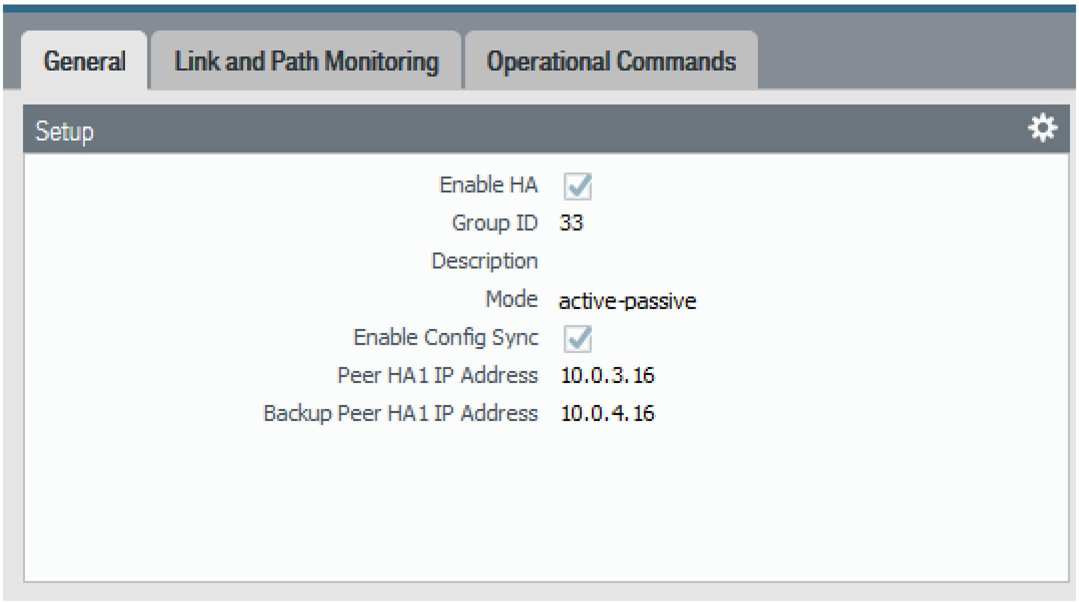

Setup

Enable HA ☑
Group ID 33
Description
Mode active-passive
Enable Config Sync ☑
Peer HA1 IP Address 10.0.3.16
Backup Peer HA1 IP Address 10.0.4.16

Enable HA
Add group ID
Choose Modes (Active-Passive)
Enable Config Sync
Set peer HA1 IP address (Port 3 IP of the Backup firewall)
Set backup peer HA1 IP address (port 4 IP of the Backup firewall)

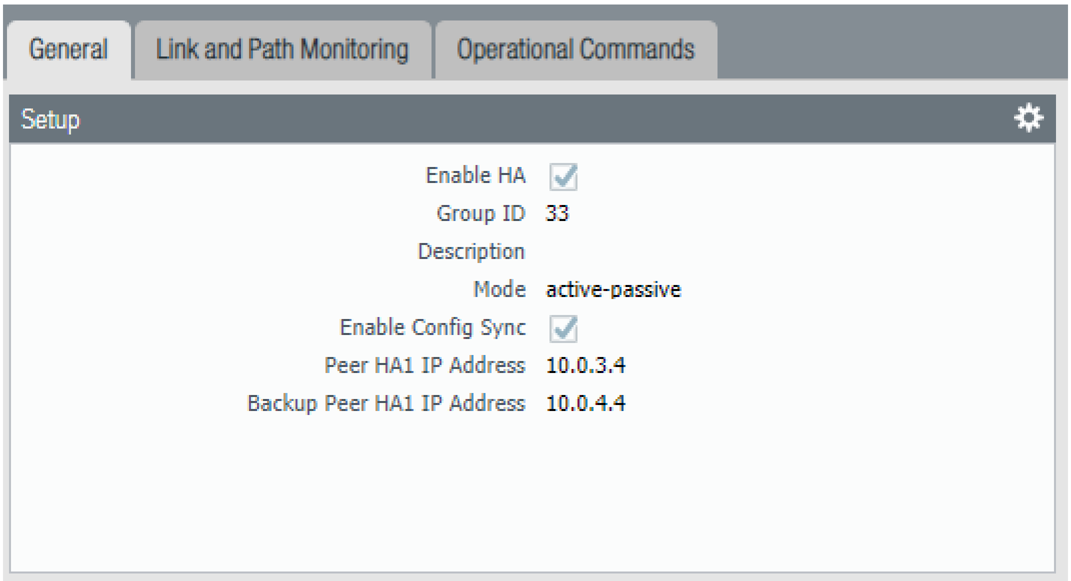

Enable HA ☑
Group ID 33
Description
Mode active-passive
Enable Config Sync ☑
Peer HA1 IP Address 10.0.3.4
Backup Peer HA1 IP Address 10.0.4.4

For the Backup Firewall Choose

Set peer HA1 IP address (Port 3 IP of the primary firewall)

Set backup peer HA1 IP address (port 4 IP of the Primary firewall)

3. From Election settings

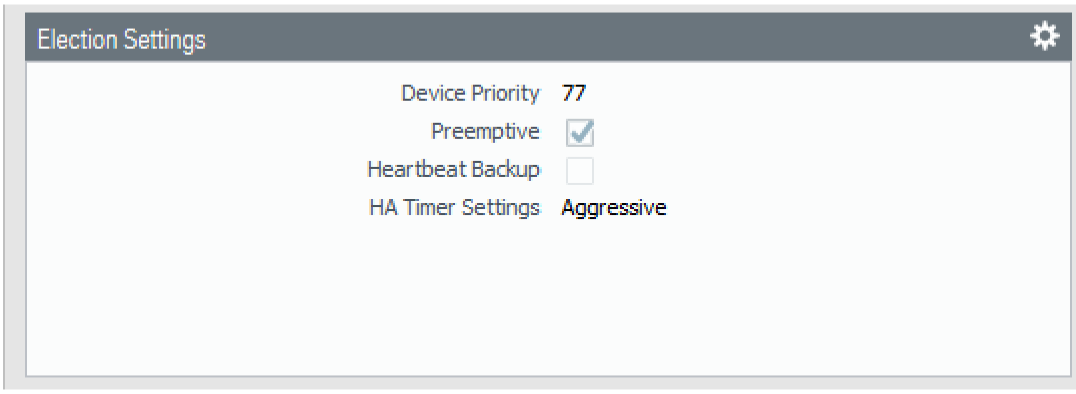

Device Priority 77
Preemptive ☑
Heartbeat Backup □
HA Timer Settings Aggressive

Set Device prioirty 77
HA Timer Settings (Aggressive)

For the Backup FW Set the device priority greater than the Active Firewall , in our case use the 78

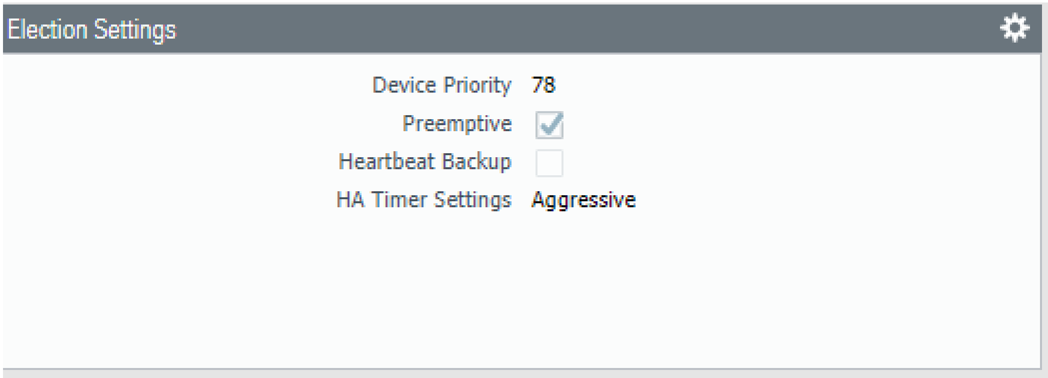

Election Settings

| | |
|---|---|
| Device Priority | 78 |
| Preemptive | ☑ |
| Heartbeat Backup | □ |
| HA Timer Settings | Aggressive |

4. From Control Link (HA1)

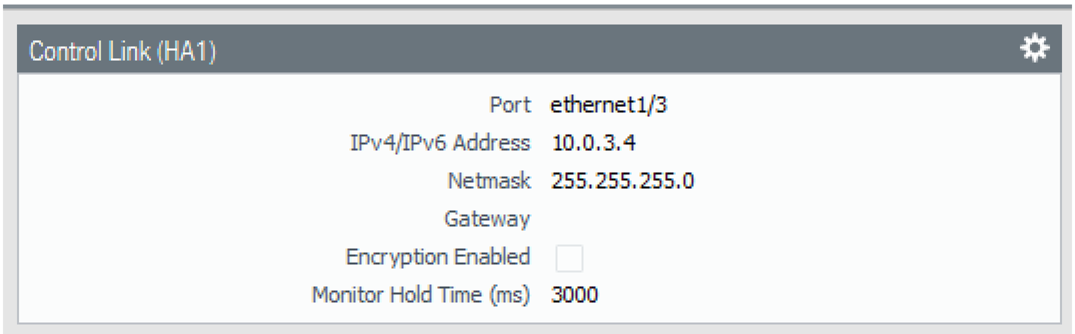

Control Link (HA1)

| | |
|---|---|
| Port | ethernet1/3 |
| IPv4/IPv6 Address | 10.0.3.4 |
| Netmask | 255.255.255.0 |
| Gateway | |
| Encryption Enabled | □ |
| Monitor Hold Time (ms) | 3000 |

Set the Port to HA port signaling/synchronizing not Data ( Port 3)

Set IP address to the HA Stnc-Subnet first port IP in the platform ( NIC 3 IP address )

For the Backup Firewall

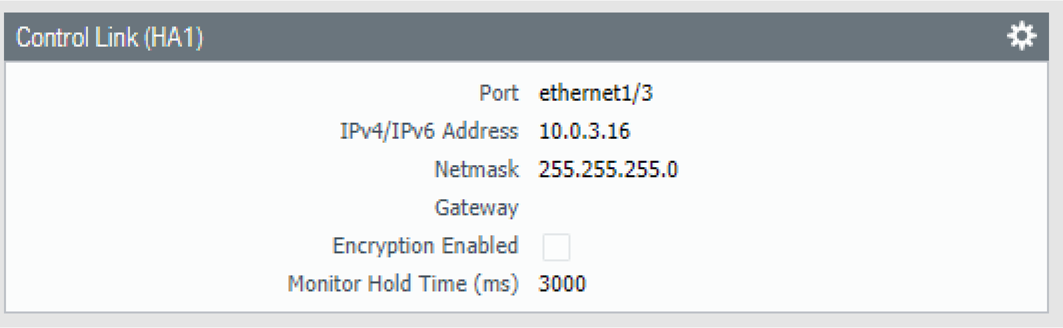

Control Link (HA1)

| | |
|---|---|
| Port | ethernet1/3 |
| IPv4/IPv6 Address | 10.0.3.16 |
| Netmask | 255.255.255.0 |
| Gateway | |
| Encryption Enabled | □ |
| Monitor Hold Time (ms) | 3000 |

Set the Port to HA port signaling/synchronizing not Data ( Port 3)

Set IP address to the HA Sync-Subnet first port IP in the platform ( NIC 3 IP address )

5. From Control Link (HA1 Backup)

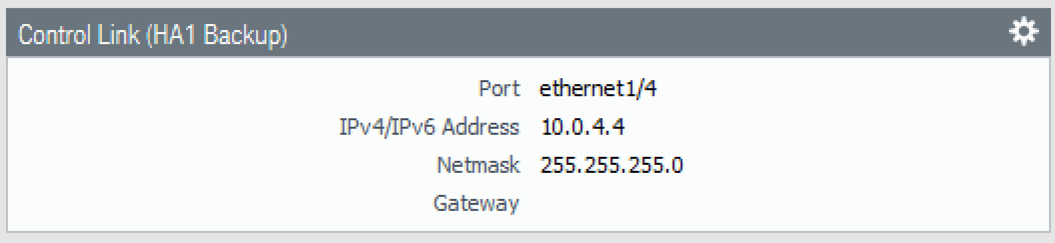

Control Link (HA1 Backup)

| | |
|---|---|
| Port | ethernet1/4 |
| IPv4/IPv6 Address | 10.0.4.4 |
| Netmask | 255.255.255.0 |
| Gateway | |

Set the Port to HA port Data ( Port 4)

Set IP address to the HA Sync2-Subnet  port IP in the platform ( NIC 4 IP address )

For the Backup Firewall

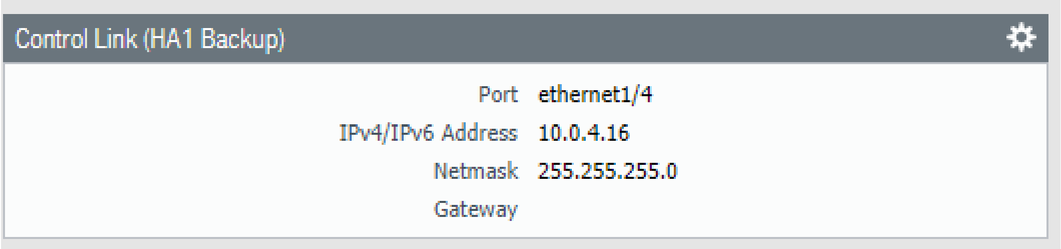

Control Link (HA1 Backup)

| | |
|---|---|
| Port | ethernet1/4 |
| IPv4/IPv6 Address | 10.0.4.16 |
| Netmask | 255.255.255.0 |
| Gateway | |

Set the Port to HA port Data ( Port 4)

Set IP address to the HA Sync2-Subnet  port IP in the platform ( NIC 4 IP address )

6. From Data Link (HA2)

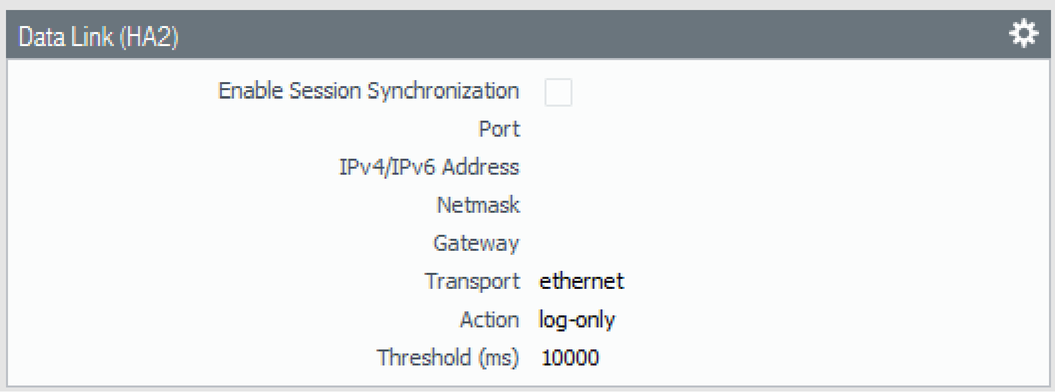

Data Link (HA2)

| | |
|---|---|
| Enable Session Synchronization | □ |
| Port | |
| IPv4/IPv6 Address | |
| Netmask | |
| Gateway | |
| Transport | ethernet |
| Action | log-only |
| Threshold (ms) | 10000 |

Perform the same high availability configuration on the backup firewall.

7. Commit

8. Restart the two firewalls. After the restart process you will find a new widget appeared on the dashboard of each firewall as shown below

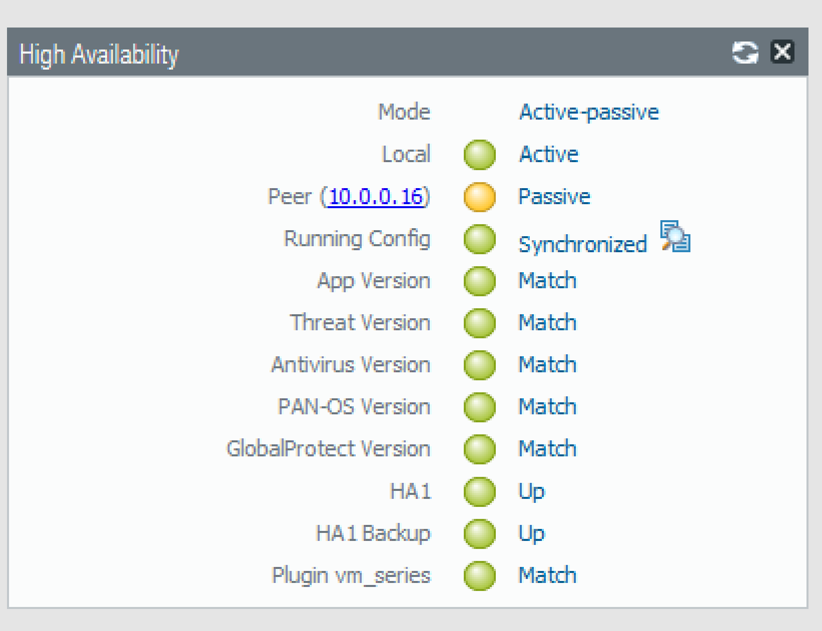

High Availability

| | | |
|---|---|---|
| Mode | | Active-passive |
| Local | ● | Active |
| Peer (10.0.0.16) | ● | Passive |
| Running Config | ● | Synchronized |
| App Version | ● | Match |
| Threat Version | ● | Match |
| Antivirus Version | ● | Match |
| PAN-OS Version | ● | Match |
| GlobalProtect Version | ● | Match |
| HA1 | ● | Up |
| HA1 Backup | ● | Up |
| Plugin vm_series | ● | Match |

10. Click on Synchronize config to transfer the initial configuration of the primary firewall to the backup firewall so that they will have identical configuration.

Primary firewall

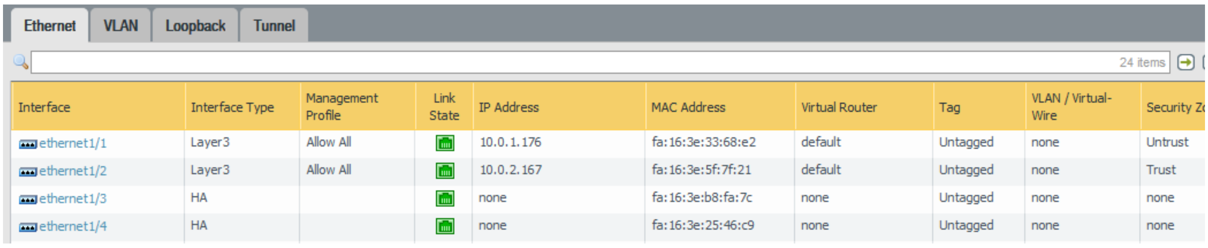

| Interface | Interface Type | Management Profile | Link State | IP Address | MAC Address | Virtual Router | Tag | VLAN / Virtual-Wire | Security Z |
|---|---|---|---|---|---|---|---|---|---|
| ethernet1/1 | Layer3 | Allow All | 🟩 | 10.0.1.176 | fa:16:3e:33:68:e2 | default | Untagged | none | Untrust |
| ethernet1/2 | Layer3 | Allow All | 🟩 | 10.0.2.167 | fa:16:3e:5f:7f:21 | default | Untagged | none | Trust |
| ethernet1/3 | HA | | 🟩 | none | fa:16:3e:b8:fa:7c | none | Untagged | none | none |
| ethernet1/4 | HA | | 🟩 | none | fa:16:3e:25:46:c9 | none | Untagged | none | none |

Backup firewall

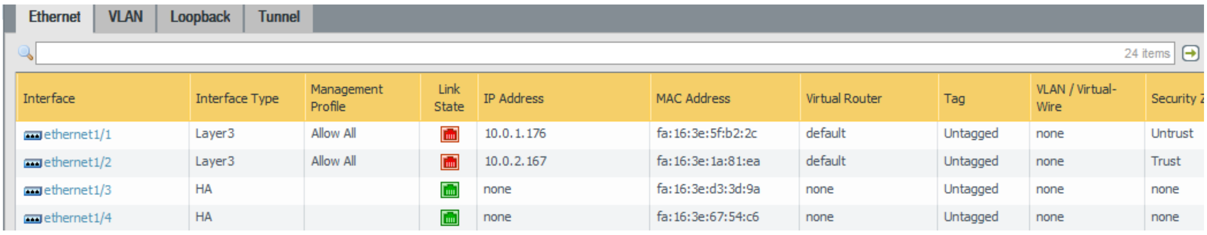

| Interface | Interface Type | Management Profile | Link State | IP Address | MAC Address | Virtual Router | Tag | VLAN / Virtual-Wire | Security Z |
|---|---|---|---|---|---|---|---|---|---|
| ethernet1/1 | Layer3 | Allow All | 🟥 | 10.0.1.176 | fa:16:3e:5f:b2:2c | default | Untagged | none | Untrust |
| ethernet1/2 | Layer3 | Allow All | 🟥 | 10.0.2.167 | fa:16:3e:1a:81:ea | default | Untagged | none | Trust |
| ethernet1/3 | HA | | 🟩 | none | fa:16:3e:d3:3d:9a | none | Untagged | none | none |
| ethernet1/4 | HA | | 🟩 | none | fa:16:3e:67:54:c6 | none | Untagged | none | none |

11. Create a Route table for the Servers Subnet ( internal subnet ) where you will configure a static route to push traffic ( incoming and outgoing ) to go to the firewall VIP  as below

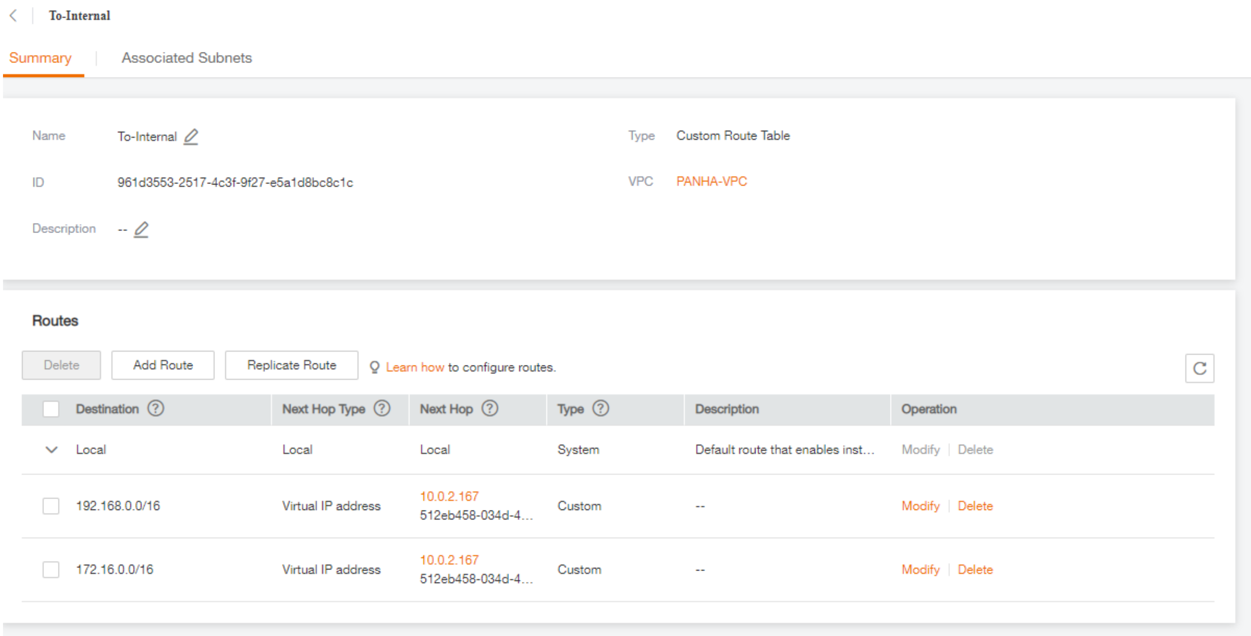

|  |  |  |  |
|---|---|---|---|
| Name | To-Internal | Type | Custom Route Table |
| ID | 961d3553-2517-4c3f-9f27-e5a1d8bc8c1c | VPC | PANHA-VPC |
| Description | -- |  |  |

Routes

| Destination | Next Hop Type | Next Hop | Type | Description | Operation |
|---|---|---|---|---|---|
| Local | Local | Local | System | Default route that enables inst... | Modify \| Delete |
| 192.168.0.0/16 | Virtual IP address | 10.0.2.167<br>512eb458-034d-4... | Custom | -- | Modify \| Delete |
| 172.16.0.0/16 | Virtual IP address | 10.0.2.167<br>512eb458-034d-4... | Custom | -- | Modify \| Delete |

# 9 IPSEC Tunnel with the on premisis Palo Alto VM Series

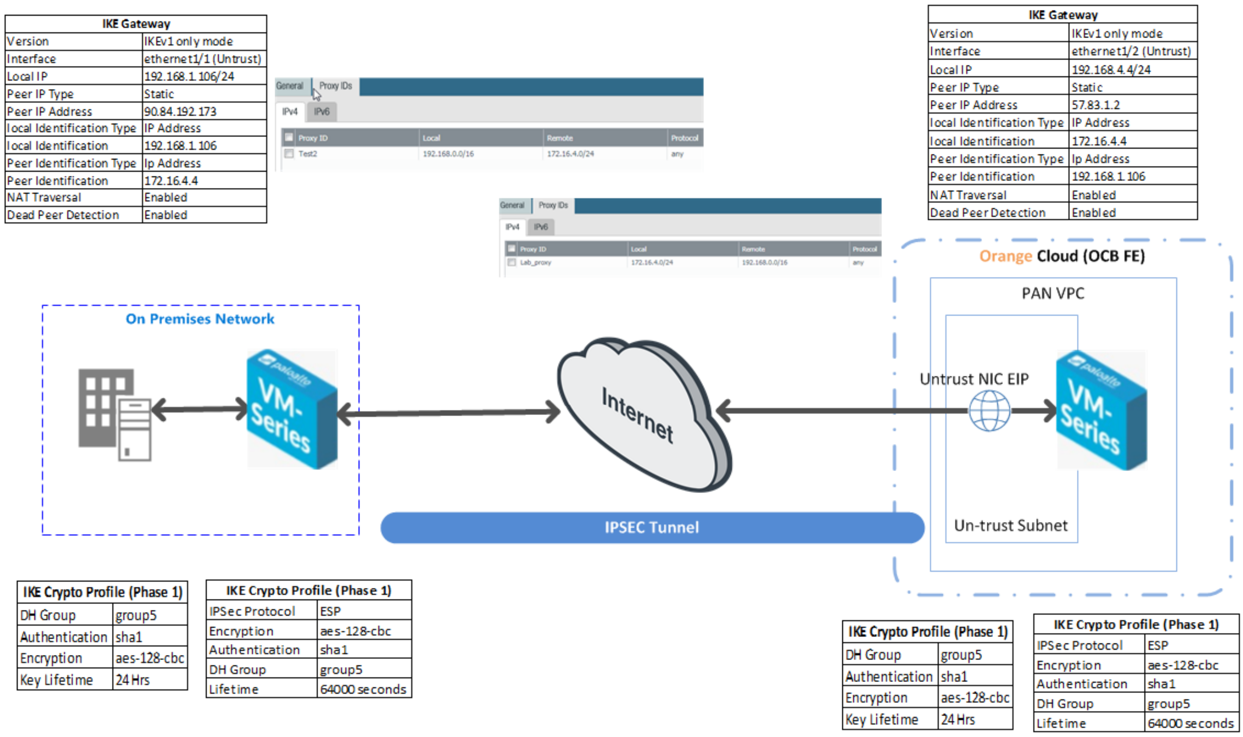

| IKE Gateway | |
|---|---|
| Version | IKEv1 only mode |
| Interface | ethernet1/1 (Untrust) |
| Local IP | 192.168.1.106/24 |
| Peer IP Type | Static |
| Peer IP Address | 90.84.192.173 |
| local Identification Type | IP Address |
| local Identification | 192.168.1.106 |
| Peer Identification Type | Ip Address |
| Peer Identification | 172.16.4.4 |
| NAT Traversal | Enabled |
| Dead Peer Detection | Enabled |

General | Proxy IDs
IPv4 | IPv6

| | Proxy ID | Local | Remote | Protocol |
|---|---|---|---|---|
| □ | Test2 | 192.168.0.0/16 | 172.16.4.0/24 | any |

General | Proxy IDs
IPv4 | IPv6

| | Proxy ID | Local | Remote | Protocol |
|---|---|---|---|---|
| □ | Lab_proxy | 172.16.4.0/24 | 192.168.0.0/16 | any |

| IKE Gateway | |
|---|---|
| Version | IKEv1 only mode |
| Interface | ethernet1/2 (Untrust) |
| Local IP | 192.168.4.4/24 |
| Peer IP Type | Static |
| Peer IP Address | 57.83.1.2 |
| local Identification Type | IP Address |
| local Identification | 172.16.4.4 |
| Peer Identification Type | Ip Address |
| Peer Identification | 192.168.1.106 |
| NAT Traversal | Enabled |
| Dead Peer Detection | Enabled |

| IKE Crypto Profile (Phase 1) | |
|---|---|
| DH Group | group5 |
| Authentication | sha1 |
| Encryption | aes-128-cbc |
| Key Lifetime | 24 Hrs |

| IKE Crypto Profile (Phase 1) | |
|---|---|
| IPSec Protocol | ESP |
| Encryption | aes-128-cbc |
| Authentication | sha1 |
| DH Group | group5 |
| Lifetime | 64000 seconds |

| IKE Crypto Profile (Phase 1) | |
|---|---|
| DH Group | group5 |
| Authentication | sha1 |
| Encryption | aes-128-cbc |
| Key Lifetime | 24 Hrs |

| IKE Crypto Profile (Phase 1) | |
|---|---|
| IPSec Protocol | ESP |
| Encryption | aes-128-cbc |
| Authentication | sha1 |
| DH Group | group5 |
| Lifetime | 64000 seconds |

IPSec Tunnel configuration will be performed on Both the firewalls as per the diagram above,

## Set Up an IPSec Tunnel

The IPSec tunnel configuration allows you to authenticate and/or encrypt the data (IP packet) as it traverses across the tunnel.
If you are setting up the Palo Alto Networks firewall to work with a peer that supports policy-based VPN, you must define Proxy IDs. Devices that support policy-based VPN use specific security rules/policies or access-lists (source addresses, destination addresses and ports) for permitting interesting traffic through an IPSec tunnel. These rules are referenced during quick mode/IKE phase 2 negotiation, and are exchanged as Proxy-IDs in the first or the second message of the process. So, if you are configuring the Palo Alto Networks firewall to work with a policy-based VPN peer, for a successful phase 2 negotiation you must define the Proxy-ID so that the setting on both peers is identical. If the Proxy-ID is not configured, because the Palo Alto Networks firewall supports route-based VPN, the default values used as Proxy-ID are source ip: 0.0.0.0/0, destination ip: 0.0.0.0/0 and application: any; and when these values are exchanged with the peer, it results in a failure to set up the VPN connection.

## Steps

1. Select Network>IPSec Tunnels and then Add a new tunnel configuration.
2. On the General tab, enter a Name for the new tunnel.
3. Select the Tunnel interface that will be used to set up the IPSec tunnel.

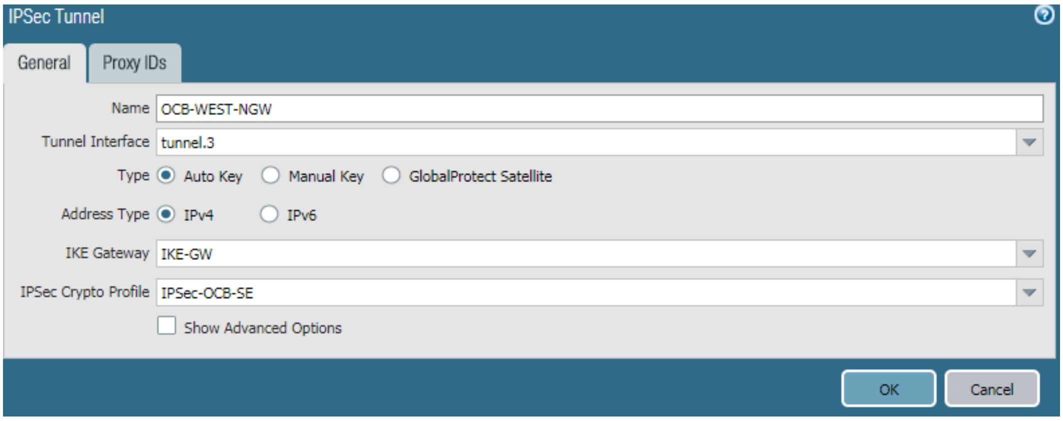

Name: OCB-WEST-NGW

Tunnel Interface: tunnel.3

Type: ● Auto Key ○ Manual Key ○ GlobalProtect Satellite

Address Type: ● IPv4 ○ IPv6

IKE Gateway: IKE-GW

IPSec Crypto Profile: IPSec-OCB-SE

□ Show Advanced Options

To create a new tunnel interface:

- Select Tunnel Interface>New Tunnel Interface. (You can also select NetworkInterfaces>Tunnel and click Add.)
- In the Interface Name field, specify a numeric suffix, such as .2.

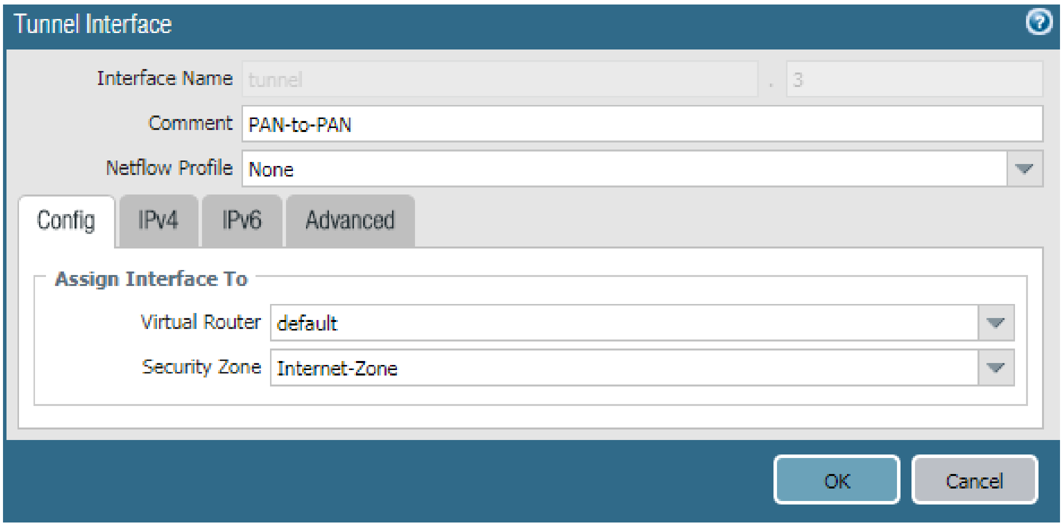

Interface Name: tunnel . 3

Comment: PAN-to-PAN

Netflow Profile: None

Virtual Router: default

Security Zone: Internet-Zone

- On the Config tab, select the Security Zone drop-down to define the zone as follows:

Use your trust zone as the termination point for the tunnel—Select the zone from the drop-down. Associating the tunnel interface with the same zone (and virtual router) as the external-facing interface on which the packets enter the firewall mitigates the need to create inter-zone routing.
Or:
Create a separate zone for VPN tunnel termination (Recommended)—Select New Zone, define a Name for the new zone (for example vpn-corp), and click OK.

- In the Virtual Router drop-down, select default.
- (Optional) If you want to assign an IPv4 address to the tunnel interface, select the IPv4 tab, and Add the IP address and network mask, for example 10.31.32.1/32.
- Click OK.

4. Define the IKE Gateway .

- Select NetworkNetwork ProfilesIKE Gateways, click Add, and on the General tab, enter the Name of the gateway.
- For Version, select IKEv1 only mode, IKEv2 only mode, or IKEv2 preferred mode. The IKE gateway begins its negotiation with its peer in the mode specified here. If you select IKEv2 preferred mode, the two peers will use IKEv2 if the remote peer supports it; otherwise they will use IKEv1. The Version selection also determines which options are available on the Advanced Options tab.

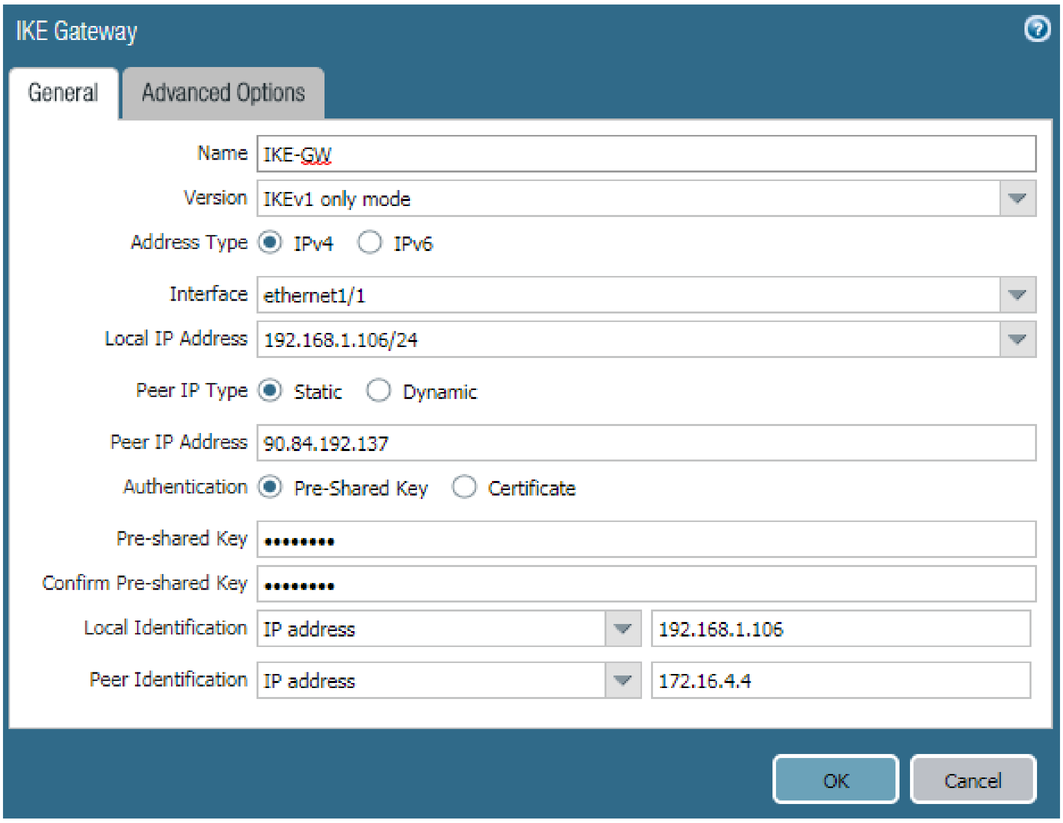

Name: IKE-GW
Version: IKEv1 only mode
Address Type: ● IPv4 ○ IPv6
Interface: ethernet1/1
Local IP Address: 192.168.1.106/24
Peer IP Type: ● Static ○ Dynamic
Peer IP Address: 90.84.192.137
Authentication: ● Pre-Shared Key ○ Certificate
Pre-shared Key: ••••••••
Confirm Pre-shared Key: ••••••••
Local Identification: IP address 192.168.1.106
Peer Identification: IP address 172.16.4.4

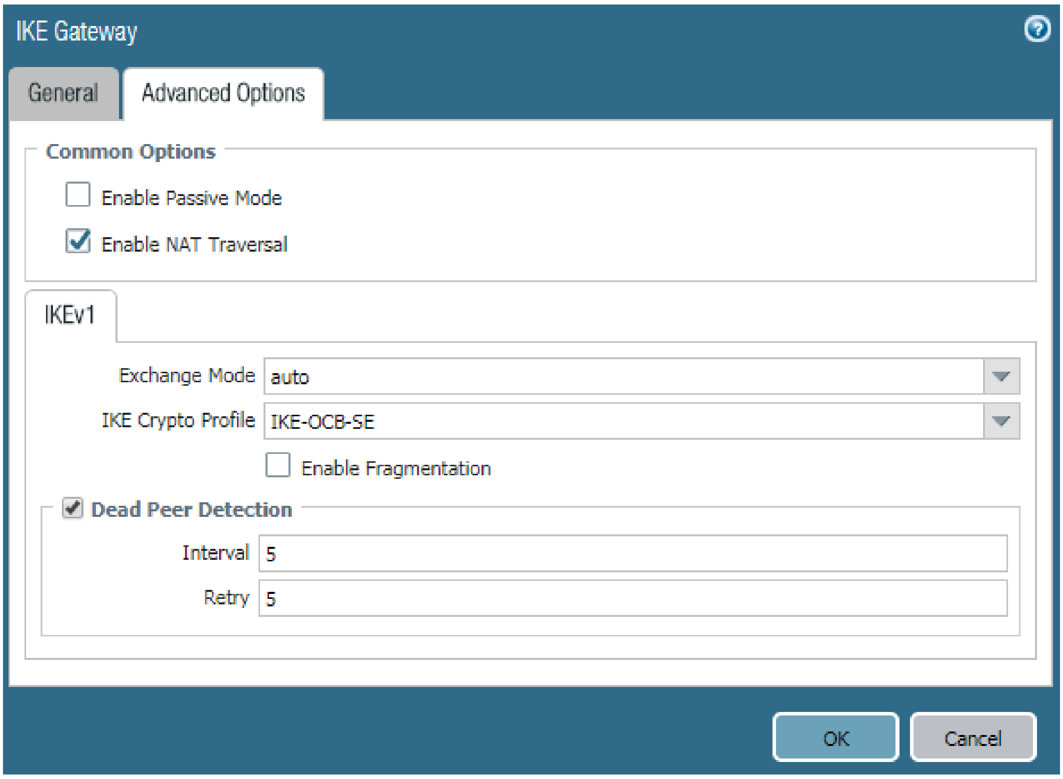

Common Options

□ Enable Passive Mode

☑ Enable NAT Traversal

IKEv1

Exchange Mode: auto

IKE Crypto Profile: IKE-OCB-SE

□ Enable Fragmentation

☑ Dead Peer Detection

Interval: 5

Retry: 5

5- Define IKE Crypto Profile

In this phase, the firewalls use the parameters defined in the IKE Gateway configuration and the IKE Crypto profile to authenticate each other and set up a secure control channel. IKE Phase supports the use of preshared keys or digital certificates (which use public key infrastructure, PKI) for mutual authentication of the VPN peers. Preshared keys are a simple solution for securing smaller networks because they do not require the support of a PKI infrastructure. Digital certificates can be more convenient for larger networks or implementations that require stronger authentication security.

When using certificates, make sure that the CA issuing the certificate is trusted by both gateway peers and that the maximum length of certificates in the certificate chain is 5 or less. With IKE fragmentation enabled, the firewall can reassemble IKE messages with up to 5 certificates in the certificate chain and successfully establish a VPN tunnel.

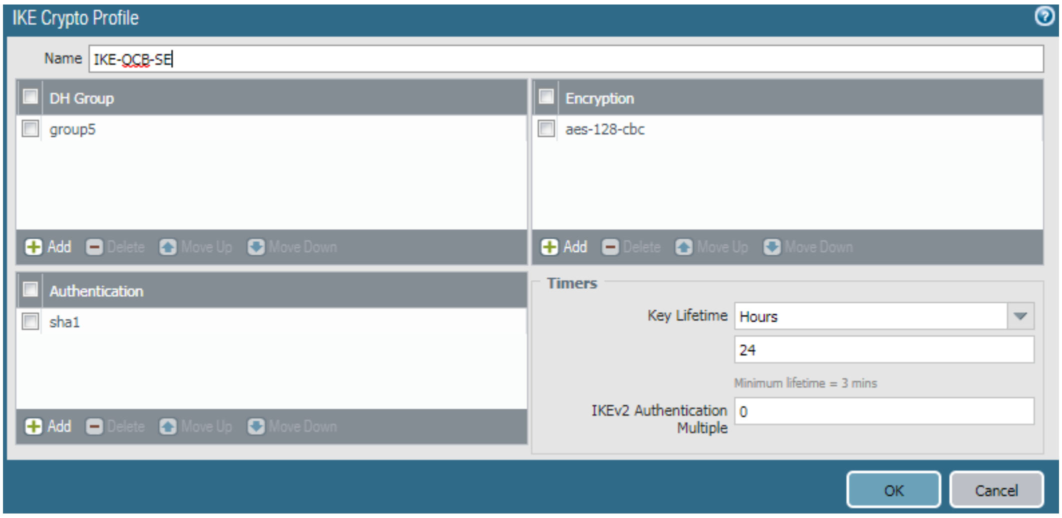

IKE Crypto Profile

Name: IKE-QCB-SE

DH Group: group5

Encryption: aes-128-cbc

Authentication: sha1

Timers

Key Lifetime: Hours

24

Minimum lifetime = 3 mins

IKEv2 Authentication Multiple: 0

6. Define IPSEC Crypto

Create a new IPSec profile.

- Select Network>Network Profiles>IPSec Crypto and select Add.
- Enter a Name for the new profile.
- Select the IPSec Protocol—ESP or AH—that you want to apply to secure the data as it traverses across the tunnel.
- Click Add and select the Authentication and Encryption algorithms for ESP, and Authentication algorithms for AH, so that the IKE peers can negotiate the keys for the secure transfer of data across the tunnel.
- Commit your IPSec profile.
- Click OK and click Commit.
- Attach the IPSec Profile to an IPSec tunnel configuration.

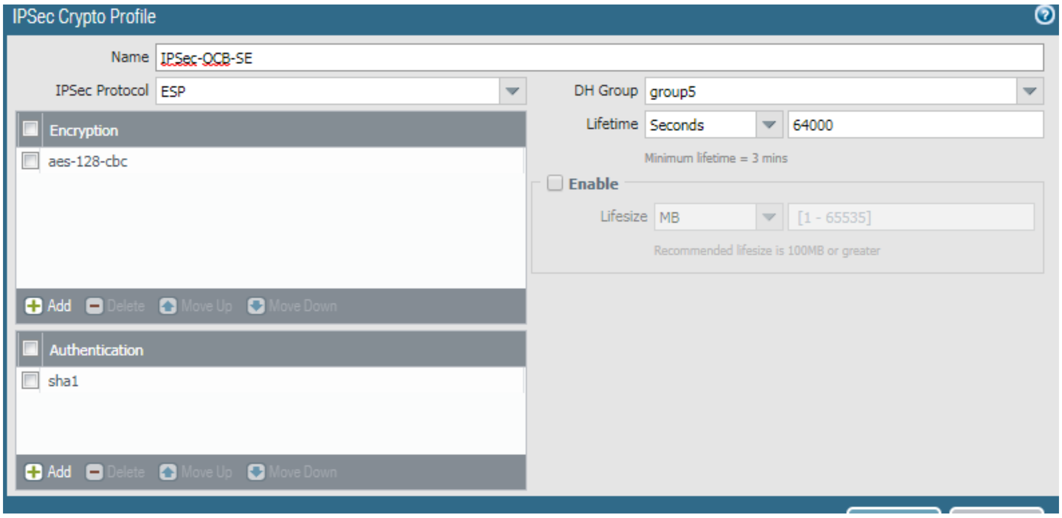

IPSec Crypto Profile

Name: IPSec-QCB-SE

IPSec Protocol: ESP

DH Group: group5

Lifetime: Seconds 64000

Minimum lifetime = 3 mins

□ Enable

Lifesize: MB [1 - 65535]

Recommended lifesize is 100MB or greater

Encryption: aes-128-cbc

Authentication: sha1

7. Setup Tunnel Monitoring (Optional)

To provide uninterrupted VPN service, you can use the Dead Peer Detection capability along with the tunnel monitoring capability on the firewall. You can also monitor the status of the tunnel. These monitoring tasks are described in the following sections:

- Define a Tunnel Monitoring Profile

A tunnel monitoring profile allows you to verify connectivity between the VPN peers; you can configure the tunnel interface to ping a destination IP address at a specified interval and specify the action if the communication across the tunnel is broken.

a. Select Network>Network Profiles>Monitor. A default tunnel monitoring profile is available for use.
b. Click Add, and enter a Name for the profile.
c. Select the Action to take if the destination IP address is unreachable.

 - o Wait Recover—the firewall waits for the tunnel to recover. It continues to use the tunnel interface in routing decisions as if the tunnel were still active.
 - o Fail Over—forces traffic to a back-up path if one is available. The firewall disables the tunnel interface, and thereby disables any routes in the routing table that use the interface.

In either case, the firewall attempts to accelerate the recovery by negotiating new IPSec keys.

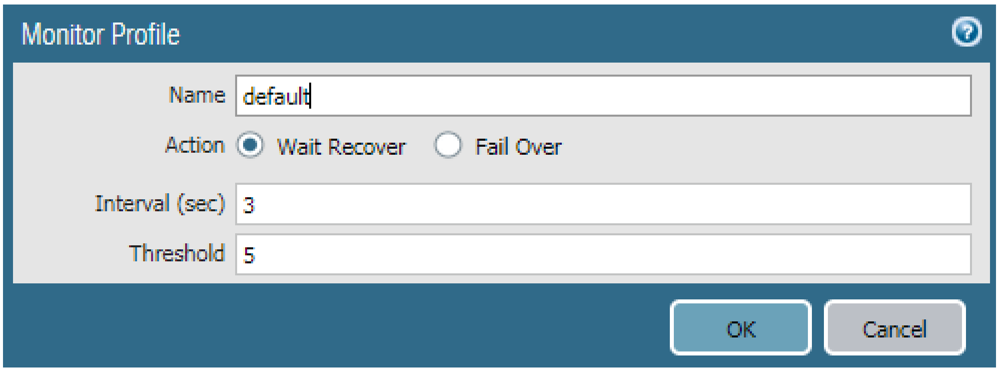

Monitor Profile

Name: default

Action: ● Wait Recover ○ Fail Over

Interval (sec): 3

Threshold: 5

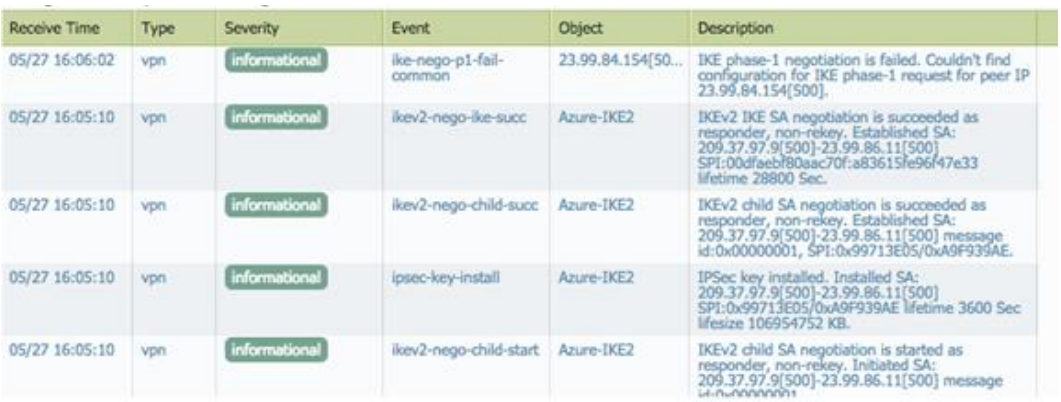

| Receive Time | Type | Severity | Event | Object | Description |
|---|---|---|---|---|---|
| 05/27 16:06:02 | vpn | informational | ike-nego-p1-fail-common | 23.99.84.154[50... | IKE phase-1 negotiation is failed. Couldn't find configuration for IKE phase-1 request for peer IP 23.99.84.154[500]. |
| 05/27 16:05:10 | vpn | informational | ikev2-nego-ike-succ | Azure-IKE2 | IKEv2 IKE SA negotiation is succeeded as responder, non-rekey. Established SA: 209.37.97.9[500]-23.99.86.11[500] SPI:00dfaebf80aac70f:a83615fe96f47e33 lifetime 28800 Sec. |
| 05/27 16:05:10 | vpn | informational | ikev2-nego-child-succ | Azure-IKE2 | IKEv2 child SA negotiation is succeeded as responder, non-rekey. Established SA: 209.37.97.9[500]-23.99.86.11[500] message id:0x00000001, SPI:0x99713E05/0xA9F939AE. |
| 05/27 16:05:10 | vpn | informational | ipsec-key-install | Azure-IKE2 | IPSec key installed. Installed SA: 209.37.97.9[500]-23.99.86.11[500] SPI:0x99713E05/0xA9F939AE lifetime 3600 Sec lifesize 106954752 KB. |
| 05/27 16:05:10 | vpn | informational | ikev2-nego-child-start | Azure-IKE2 | IKEv2 child SA negotiation is started as responder, non-rekey. Initiated SA: 209.37.97.9[500]-23.99.86.11[500] message id:0x00000001 |

8. From Network > IPSec Tunnels > Add new

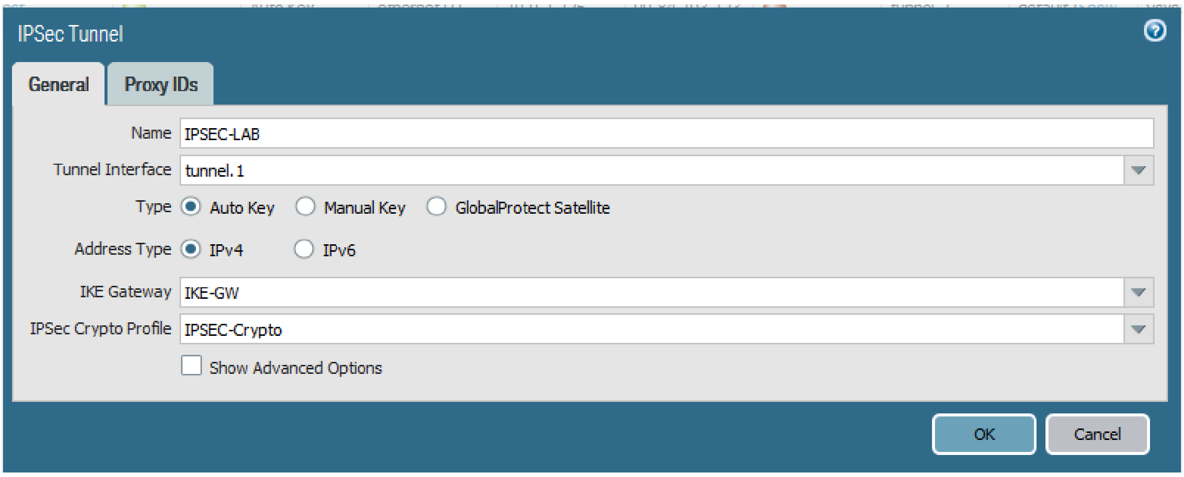

IPSec Tunnel

| | |
|---|---|
| Name | IPSEC-LAB |
| Tunnel Interface | tunnel.1 |
| Type | ● Auto Key ○ Manual Key ○ GlobalProtect Satellite |
| Address Type | ● IPv4 ○ IPv6 |
| IKE Gateway | IKE-GW |
| IPSec Crypto Profile | IPSEC-Crypto |
| | □ Show Advanced Options |

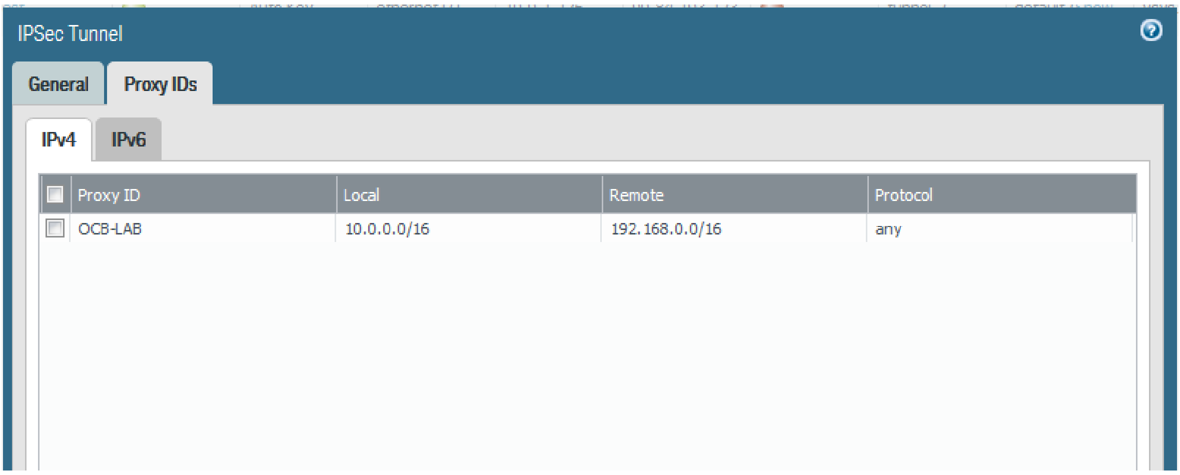

IPSec Tunnel

| □ | Proxy ID | Local | Remote | Protocol |
|---|---|---|---|---|
| □ | OCB-LAB | 10.0.0.0/16 | 192.168.0.0/16 | any |

9. Now the Ipsec tunnel is set between the OCBFE and on premisis network

OCB FE Primary Firewall

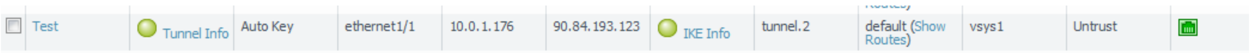

| | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| □ | Test | 🟢 Tunnel Info | Auto Key | ethernet1/1 | 10.0.1.176 | 90.84.193.123 | 🟢 IKE Info | tunnel.2 | default (Show Routes) | vsys1 | Untrust | 🟩 |

OCB FE Backup Firewall

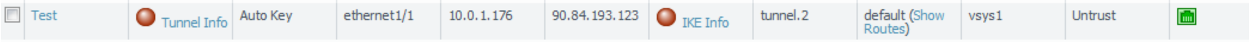

| | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| □ | Test | 🔴 Tunnel Info | Auto Key | ethernet1/1 | 10.0.1.176 | 90.84.193.123 | 🔴 IKE Info | tunnel.2 | default (Show Routes) | vsys1 | Untrust | 🟩 |

On premisis Firewall

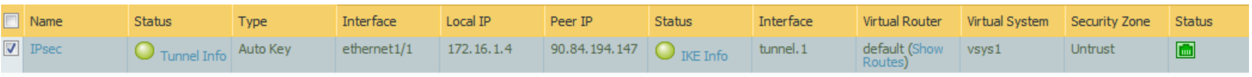

| □ | Name | Status | Type | Interface | Local IP | Peer IP | Status | Interface | Virtual Router | Virtual System | Security Zone | Status |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| ☑ | IPsec | 🟢 Tunnel Info | Auto Key | ethernet1/1 | 172.16.1.4 | 90.84.194.147 | 🟢 IKE Info | tunnel.1 | default (Show Routes) | vsys1 | Untrust | 🟩 |