



PAN-OS Active/Active High Availability

Configuring active/active clusters

May 2011

Jerish Parapurath
Palo Alto Networks
232 E. Java Dr.
Sunnyvale, CA 94089
408.738.7700
www.paloaltonetworks.com

Table of Contents

Overview.....	3
Hardware requirements.....	3
Software requirements.....	3
Feature description.....	3
Packet handling within Active-Active cluster.....	4
Session Ownership.....	4
Session Setup.....	4
Packet flow in a cluster.....	5
New session.....	5
Packet matching an existing session.....	6
Asymmetric Flow.....	6
Active-Active deployment options.....	7
Virtual Wire mode.....	7
Layer 3 mode.....	8
Floating IP.....	8
ARP load sharing.....	9
Mixed mode: Combining Floating IP and ARP load sharing.....	9
Route based redundancy.....	10
Virtual MAC.....	11
Configuring Active-Active cluster.....	12
Active-Active device states.....	12
Connecting the devices.....	12
HA links.....	13
Setting the device mode.....	13
Election settings.....	15
Control link.....	16
Data link.....	16
Backup HA links.....	17
HA3 link.....	17
Session setup options.....	18
Configuring virtual addresses.....	19
Case1: Using floating IP and ARP load sharing.....	19
Configuring shared IP- ARP load sharing.....	20
Configuring Floating IP.....	22
Case2: Using floating IP.....	22
Configuring NAT with active-active cluster.....	23
Source NAT using interface IP.....	24
Verification.....	25
Source NAT using IP pools.....	26
Source NAT rules with device failure.....	27
Destination NAT in layer2 connected network.....	27
Destination NAT in layer3 connected network.....	29
IPSec with active-active cluster.....	30
Configuration.....	30
Tunnel status.....	31
SSL VPN.....	32
Verifying SSL flow.....	33
Logs and packet capture.....	33
Summary.....	33

Overview

Service providers and enterprises that deliver revenue-generating and business critical services over the Internet face a myriad of performance and security challenges. However critical those challenges may be, high availability remains the paramount concern. In order to properly perform access control functions, a network firewall must be placed at the single point through which all traffic must pass. Because all traffic must pass through the firewall, it is vital that the traffic flow remains uninterrupted, even in the event of a device or network failure. A well-designed security infrastructure needs to offer high availability tools to create a resilient, scalable, and easy to manage solution.

Palo Alto Networks offers a line of purpose-built security solutions that integrate firewall and VPN functions with a set of high availability (HA) tools to deliver resilient, high performance devices. This technical paper describes the main functionality of PAN-OS high availability

Hardware requirements

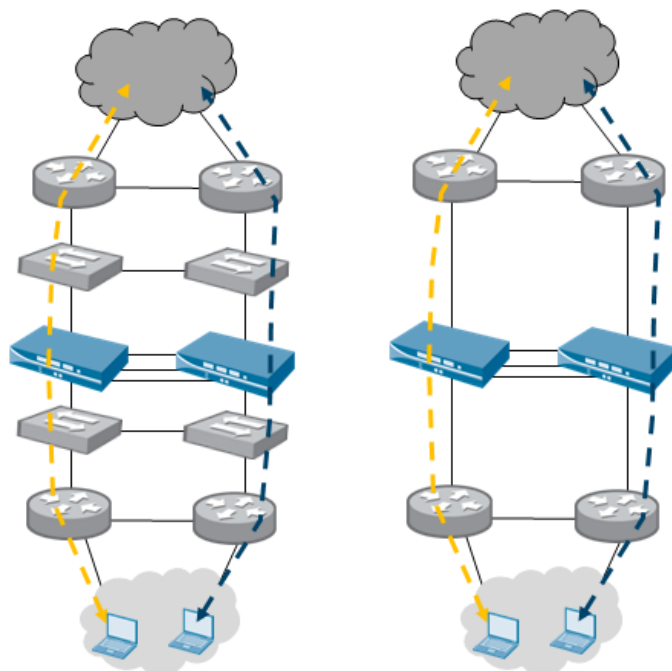
Two identically configured Palo Alto Networks next-generation firewalls per cluster.

Software requirements

PAN-OS 4.0 and later.

Feature description

A PAN-OS HA cluster consists of two identical Palo Alto Networks next-generation firewalls with identical software that enforce the same overall security policy and share the same configuration settings. With Active-Active deployment, both the devices are active and processing traffic. Active-Active HA is supported only in the virtual-wire and Layer 3 modes. Such deployments are most suited for scenarios involving asymmetric routing. In addition to the HA1 and HA2 links used in active-passive, active-active deployments require a dedicated HA3 link. This link is used as packet forwarding link for session setup and asymmetric traffic handling.



Packet handling within Active-Active cluster

In an active-active cluster, the packet handling can be distributed between the two devices. There are two important functions that are handled by the devices in a cluster

- Session ownership
- Session setup

Session Ownership

Within an active-active cluster, the session owner device can be either the firewall that receives the first packet of a new session or the device in an active-primary state. This device is responsible for all layer 7 processing, i.e. app-id, content-id, and threat scanning for this session. This device is also responsible for generating all traffic logs for the session.

Session Setup

The session setup device is responsible for the layer2 through layer4 processing required for setting up a new session. Address translation is performed by the session setup device. The session setup device is determined by configuring the “session setup load sharing” options. The separation of session owner and session setup devices is necessary to avoid race conditions that can occur in asymmetrically routed environments

Note:

- All packet forwarding between the two devices uses a dedicated link, also known as the HA3 link
- When the session owner fails, the peer device will become the session owner. The existing sessions will fail over to the functional device and no layer 7 processing will be available for these sessions.


- When a device recovers from a failure, all sessions that were owned by the device before failure will revert back to the original device

Packet flow in a cluster

In order to understand the packet flow within a cluster, we will discuss three different scenarios

1. New session
2. Established session
3. Asymmetric packet flow


New session



The sequence of steps involved in setting up a session is listed below



1. End host sends packet to device-A.
2. Firewall examines the contents of the packet to match it to an existing session.
3. If no session match, Dev-A determines that it has received the first packet for a new session. Therefore Dev-A becomes the session owner.
4. Dev-A uses the configured session setup load sharing options to identify the session setup device. In this example we assume the setup function is performed by Dev-B
5. Using the HA-3 link, Dev-A sends the first packet it received to Dev-B.
6. Dev-B sets up the session and returns the packet to Dev-A for layer 7 processing if any.
7. Dev-A then forwards the packet out via the egress interface to the destination

Packet matching an existing session



1. End host sends packet to Dev-A.
2. Firewall examines the contents of the packet to match the packet to an existing session
3. If there is a session match, Dev-A processes the packet and sends the packet out via the egress interface to the destination

Asymmetric Flow



1. Dev-B receives a packet.
2. Receiving device has a session for the packet but it is owned by peer device, Dev-A.

3. Dev-B forwards packet over the HA3 link to the Dev-A for processing.
4. In Vwire deployment in order to preserve the forwarding path, Dev-A processes the packet and returns to Dev-B, to be transmitted out the egress interface to the destination.
5. In layer3 deployment , Dev-A processes packet and forwards it to destination if it has the route



Active-Active deployment options

Active-Active HA is supported only in Virtual-Wire and Layer 3 modes. In the following sections we will discuss different supported deployment modes of active-active cluster

Virtual Wire mode

Active- Active HA in Vwire mode offers the simplest solution to implement high availability. As shown in the figure below, the firewalls are installed between L3 devices. These are often used in conjunction with dynamic routing protocols which will fail traffic over to the other cluster member if needed.

Note: Implementing A/A HA in vwire mode in a layer2 sandwich will result in switching loops if Spanning Tree Protocol is not enabled on the switches. It is recommended to deploy A/A in vwire in a layer3 topology.




Layer 3 mode

Layer3 mode of deploying active-active HA, supports the use of virtual IP addressing, NAT, and the use of dynamic routing protocols for redundancy. An Active-active cluster can be deployed in several different scenarios in layer3 mode as described below

Floating IP

This deployment option allows for the creation of floating IP addresses that can move between the HA devices when a link failure or device failure occurs. The interface on the device in the cluster that owns the floating IP address responds to ARP requests with a virtual MAC address. Floating IP addresses are recommended when VRRP- like functionality is required. Floating IP addresses can also be used to implement VPNs and source Network Address Translation (NAT) configurations, allowing for persistent connections when a failure occurs on the device offering those services. As shown in the figure below, each interface on the firewall has its own IP address and a floating IP address configured. Please note, the interface IP address remains local to the device but the floating IP address can move between the devices upon device failure. The end hosts are configured to use the floating IP address as the default gateway, allowing the traffic to be load balanced within the cluster. External load balancers can also be used to load balance traffic between with firewalls within the cluster.




In the event of a link or device failure the floating IP and VMAC moves over to the functional device. A Gratuitous ARP is sent out the by the functional device to update the MAC table of the connected switches. Once the device recovers from the failure, the floating IP and VMAC will move back to the original device.

ARP load sharing


ARP load-sharing allows the HA pair to share an IP address and provide gateway services. In this scenario, all hosts are configured with a single gateway IP address. ARP requests for the gateway IP address are responded to with a virtual MAC address from a single device in the pair. Each device will have unique virtual MAC address generated for the shared IP. The device that responds to the ARP request is determined by computing the hash or modulo of the source IP address of the ARP request. It is important to note that once the end host receives the ARP response from the gateway, it caches the MAC address and all traffic from the host is routed via the firewall that responded with the VMAC for the life time of the ARP cache. Life time of the ARP cache is dependent on the end host operating system.

ARP load-sharing should be used only when a Layer 2 separation exists between the firewall and end hosts. In the event of a link or device failure, the floating IP and VMAC moves over to the functional device. Gratuitous ARP is sent out the by the functional device to update the MAC table of the connected switches.




Mixed mode: Combining Floating IP and ARP load sharing

It is possible to have some of the interfaces configured with floating IPs and some with shared IPs for ARP load sharing. As show in the figure below, the cluster can be configured with ARP load sharing IPs, configured for the hosts on the LAN segment, and floating IP address configured on the upstream WAN edge routers.



Route based redundancy

Dynamic routing protocols can be used to determine the best path and load share between a pair of firewalls in an active-active cluster. In this case no floating IP addresses are configured on the interface. Each device in the cluster has a unique IP address configured to each one of the interfaces. The IP address remains local to the device and does not move between devices in the event of a device failure.



Virtual MAC

In layer 3 mode a virtual MAC address is created for the all the configured floating IP address and ARP load sharing IP address interfaces. The format of the virtual MAC is 00-1B-17:00: xx: yy where

00-1B-17: vendor ID

00: fixed

xx: The format is shown below

7	6	5	4	3	2	1	0
1	Device-ID		HA group-ID				

yy: interface ID

When a new active device takes over, Gratuitous ARPs are sent from each of the connected interfaces of the functional device to inform the connected Layer-2 switches of the new location of the virtual MAC addresses

Configuring Active-Active cluster

In this section we will discuss configuring an active-active cluster in different deployment scenarios. Before getting into the configuration details, let us discuss the different device states

Active-Active device states

The following are the possible HA system states:

Initial: The transient state of a device when it joins the HA cluster. The device will remain in this state for 60 seconds after bootup unless a peer is discovered and negotiation begins. After 60 seconds, the device will become active if HA negotiation has not started.

Active-Primary: The state of the device that connects to User-ID agents, runs DHCP server/relay, matches NAT and PBF rules with device ID of primary

Active-Secondary: The active state of the device in active-active cluster. Device in this state can own sessions, and setup sessions

Tentative: State of the device caused by a monitored object failure. A device in this state will synchronize sessions, and configurations from the peer device.

In vwire mode, when a device goes into tentative state due to path monitoring failure, and receives a packet to be forwarded, it will send that traffic to the peer device via the HA3 link for processing. The peer device will process the traffic and send it back via the HA3 link to the device to be sent out via egress interface. This is done to preserve the forwarding path in vwire mode. In Layer 3 mode, when a device in tentative state receives a packet, it will send that packet over the HA3 link for the peer device to own/setup the session.

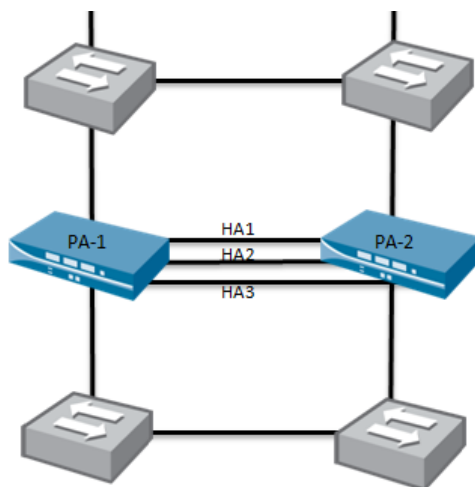
Depending on the network topology, this device will either send the packet out to the destination or send it back to the device in tentative state for forwarding.

Non-functional: Error state due to a data plane crash or a configuration mismatch.

Suspended: This is an administratively disabled state. In this stage HA cluster member cannot participate in the HA election process and does not sync configurations and sessions


Connecting the devices

The two Palo Alto Networks next-generation firewalls are cabled are shown in the figure above. It is important to note that there are three HA interfaces. These interfaces are HA1- control link, HA2- data link and HA3- packet forwarding link.



HA links

The devices in a HA cluster require two dedicated connections to exchange state information and data synchronization. The HA3 link is strictly used for packet forwarding. The figure below shows the logical representation of two devices connected in a active-active HA cluster



Setting the device mode

The first step in configuring Active-Active cluster is to set the HA mode to active-active. This also requires configuring ID, device ID and IP address of the peer device

Device>high availability>setup

Enable HA	<input checked="" type="checkbox"/>
ID	<input type="text" value="1"/> (1 - 63)
Description	<input type="text"/>
Mode	active-active ▾
Device Id	<input type="text" value="0"/> <small>Device ID in HA group, 0 or 1</small>
Peer HA IP Address	<input type="text" value="1.1.1.100"/>
Backup Peer HA IP Address	<input type="text"/>
Enable Config Sync	<input checked="" type="checkbox"/>

ID: This is the HA group ID. Both devices must have the same group ID. HA group-ID is used to calculate virtual MAC.

Mode: Choose active-active from the drop down list

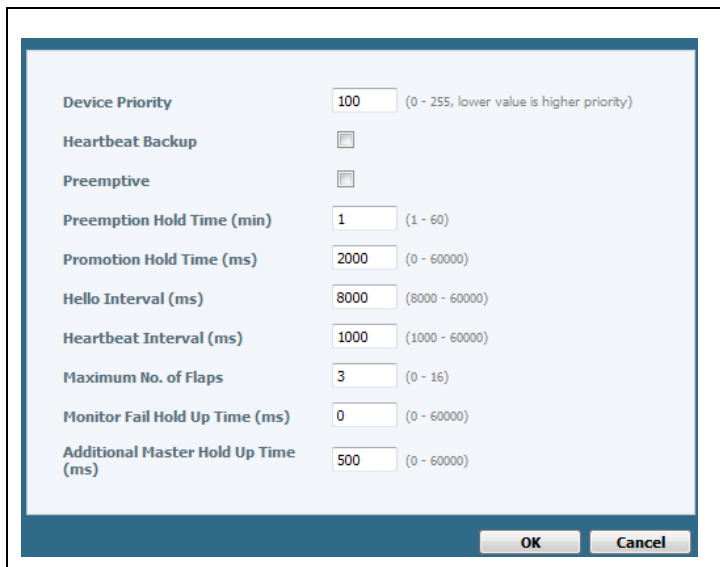
Device-id: Select unique device from the drop down list. The device-id can be either 0 or 1. The device-ID remains local to the device and does not transition between devices during failover. This field is also used to calculate virtual MAC address.

Peer HA IP Address: IP address of the HA-1 control link on the peer device

Backup Peer HA IP Address: IP address of the backup control link on the peer device. This field is optional

Enable Config Sync: This feature is enabled by default. This is required to synchronize configuration between the devices in cluster

Election settings



Device Priority	100	(0 - 255, lower value is higher priority)
Heartbeat Backup	<input type="checkbox"/>	
Preemptive	<input type="checkbox"/>	
Preemption Hold Time (min)	1	(1 - 60)
Promotion Hold Time (ms)	2000	(0 - 60000)
Hello Interval (ms)	8000	(8000 - 60000)
Heartbeat Interval (ms)	1000	(1000 - 60000)
Maximum No. of Flaps	3	(0 - 16)
Monitor Fail Hold Up Time (ms)	0	(0 - 60000)
Additional Master Hold Up Time (ms)	500	(0 - 60000)

OK Cancel

Device Priority: Enter numeric value for the device priority. Upon initial configuration the device with the lowest priority (value closest to zero) becomes the active unit (default priority is 100). If two devices have the same priority value, the device with device-id 0 becomes the active-primary unit.

Heartbeat backup: Select this to use the management ports on the firewall to provide a backup path for heartbeat and hello messages

Preemptive: Select the check box to enable the higher priority firewall (priority numeric value closest to zero) to resume active-primary operation after recovering from a failure. If this setting is off, then the lower priority firewall remains active-primary even after the higher priority firewall recovers from a failure. This option must be enabled to on both the devices for preemptive to work.

Preemption-hold-time: The time in minutes a device remains in the passive state before taking over as the active device (default 1 min). If the two devices are configured with different preemption hold-down timeouts, then the preemption hold-down timeout configured on the higher priority (lower value) device is used. The preemption hold-down timeout on the lower priority device is ignored

Promotion hold time: The time a device stays in the active-secondary state before changing state to active-primary, In the event the active-primary device has transitioned to a non-functional or tentative state because of a failure, the active-secondary device will immediately transition to an active-primary state regardless of the value defined as the promotion hold time.

Hello-interval: It is the time interval in milliseconds to send Hello messages to track the status of the HA agent. Hellos are sent over the HA1 link. The default value of the hello interval is platform dependent. The PA-500 and PA-2000 Series uses 8000ms and PA-4000 Series and the PA-5000 series use 1000ms.

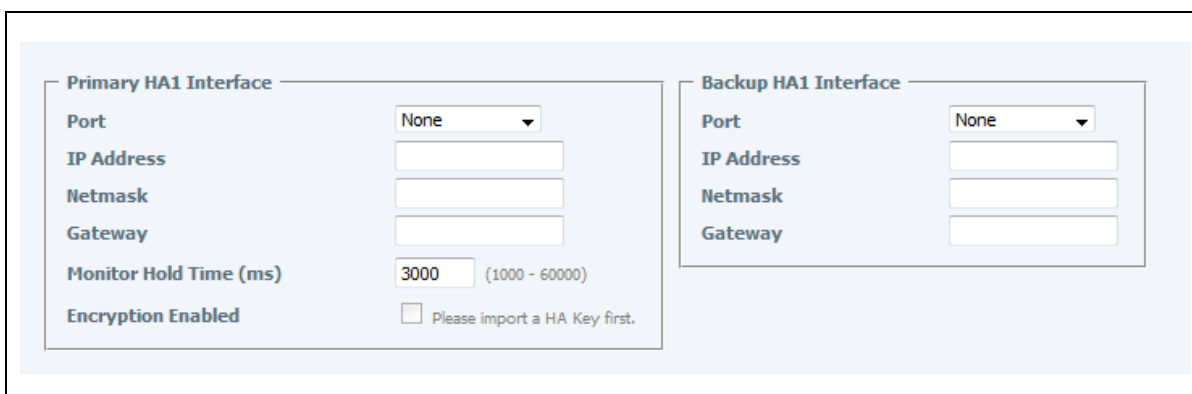
Heartbeat-interval: The time interval in milliseconds to send ICMP ping to the HA peer over the control link. The peer kernel directly responds to the pings. The default value of the heartbeat interval for all platforms is 1000ms

Maximum number of flaps: A flap is counted as the device changes state from active to non-functional. If the device changes from active to non-functional to passive to active and non-functional, 3 times within 15 minutes, the device enters suspended state. Flap max is the maximum number of HA state changes from either active to non-functional before entering suspended state.

Monitor Fail Hold up time: The time the device waits before transitioning to the tentative state because of a monitored object failure. Use this timer to prevent a device changing state to tentative due to a port flapping on the directly connected upstream device

Additional Master Hold up time: Time the ACTIVE-PRIMARY device stays ACTIVE-PRIMARY upon a monitored object failure. If the monitor hold up time is configured, the device waits MASTER HOLD UP TIME+ MONITOR HOLD UP TIME before changing state to tentative. This timer should be set if a particular device in a HA cluster is preferred as the active-primary device.

Control link



The screenshot shows the configuration page for HA1 interfaces. It is divided into two main sections: 'Primary HA1 Interface' and 'Backup HA1 Interface'. Each section contains the following fields:

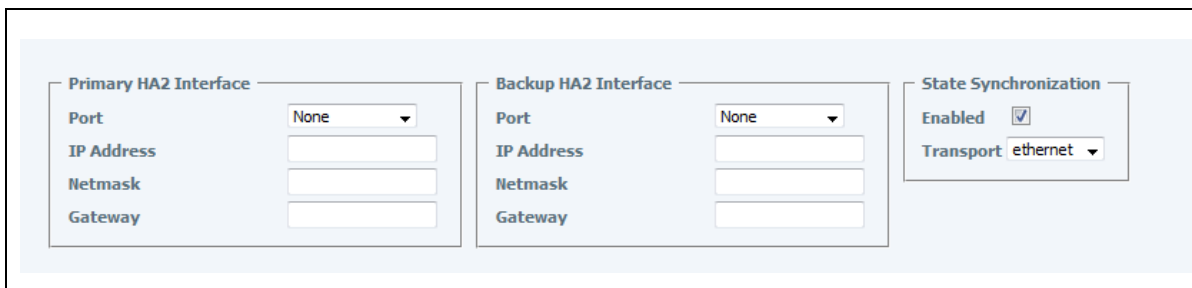
- Port:** A dropdown menu currently set to 'None'.
- IP Address:** A text input field.
- Netmask:** A text input field.
- Gateway:** A text input field.
- Monitor Hold Time (ms):** A numeric input field with the value '3000' and a range '(1000 - 60000)'.
- Encryption Enabled:** A checkbox that is currently unchecked, with the text 'Please import a HA Key first.' next to it.

The HA control link also known as the HA1 link is used by the HA agent for the devices in HA to communicate with one another. The HA1 link is a layer3 link requiring an IP address. The HA agent uses TCP port 28769 for clear text communication, or SSH over TCP port 49969 if using encryption. This connection is used to send and receive hellos and HA state information, and configuration sync and management plane sync, such as routing and user-id information. Configuration changes to either units are automatically synchronized to the other device over this link

The PA-4000 Series and PA-5000 Series firewalls have dedicated HA1 links. All other platforms require a revenue port to be configured as a HA1 link.

Monitor Hold time: HA control link monitoring tracks the state of the HA1 link to see if the peer HA device is down. This will catch a power-cycle, a reboot, or a power down of the peer device. To ignore the flapping of a link that wouldn't necessarily take the HA control connection down, a monitor hold down timer for the HA control link monitoring can be configured. The monitor hold down time is configured under the HA1 link. The default value is 3000ms.

Data link



The screenshot shows the configuration page for HA2 interfaces and state synchronization. It is divided into three main sections:

- Primary HA2 Interface:** Contains fields for Port (dropdown, 'None'), IP Address, Netmask, and Gateway.
- Backup HA2 Interface:** Contains fields for Port (dropdown, 'None'), IP Address, Netmask, and Gateway.
- State Synchronization:** Contains an 'Enabled' checkbox (checked) and a 'Transport' dropdown menu set to 'ethernet'.

The HA data link, also known as the HA2 link, is used to synchronize state, sessions, routing tables, IPSec security associations, and ARP tables between devices in a HA cluster. The HA 2 link is a layer 2 link. It uses Ethernet as transport by default with Ether Type 0x7261. Use this setting when connecting the HA2 link via switch or back-to-back. The HA data link can also be configured to use either IP or UDP as the transport. This allows the HA data link to be connected across layer3 networks. UDP encapsulation offers a benefit in that UDP computes a checksum of the entire packet, as compared to IP, where the checksum is computed only on the IP header.

IP protocol number 99 and UDP port 29281 are used when using IP or UDP as the transport. The PA-4000 and 5000 series of firewalls have dedicated HA2 link. All other platforms require a revenue port to be configured as a HA2 link.

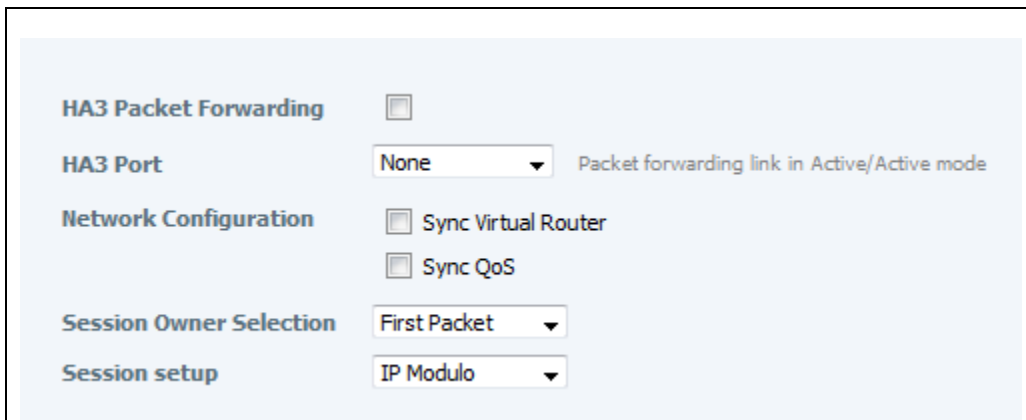
Note: The PA-4000 and 5000 series of firewalls also support using revenue ports as HA1 and HA2 links. Although it's possible to use fiber ports for HA1, it is not a recommended use. The dedicated HA1 port allows for a direct connection between the Control Planes of the two HA devices. By using an in-band port, the control messages will be following a less direct route to the control plane through the dataplane. The possibility of a software fault on the dataplane can cause a lost HA control message with adverse effects.

Backup HA links

In order to provide redundancy one of the data ports can be configured as a backup link for both HA1 and HA2 ports. The following guidelines apply when configuring backup HA links

- The IP addresses of the primary and backup HA links must not overlap each other.
- HA backup link must be in its own subnet.
- HA-1 backup and HA2-backup ports must not overlap each other. Separate links (one each) for HA1 and HA2 must be used as backup ports.

HA3 link



The screenshot shows a configuration interface for HA3 link settings. It includes the following options:

- HA3 Packet Forwarding:**
- HA3 Port:** None (dropdown menu) Packet forwarding link in Active/Active mode
- Network Configuration:**
 - Sync Virtual Router
 - Sync QoS
- Session Owner Selection:** First Packet (dropdown menu)
- Session setup:** IP Modulo (dropdown menu)

The HA3 link is used for packet forwarding between the session owner and the session setup device in an active-active cluster. HA3 link is a layer2 link and uses MAC-in-MAC encapsulation. Aggregate interfaces can be configured as a HA3 link on the PA-5000 and PA-4000 Series. This also provides redundancy of HA3 link. The interface that will be used as HA3 link must be set as type HA

```
set network interface ethernet <value> ha
```

Note: No backup links can be configured for the HA3 link. Use the aggregate interface on the PA-5000 and PA-4000 series if redundancy is required.

Session Owner options

The session ownership is user configurable. There are two options available

- First packet
The device that receives the first packet from the end host will own the session
- Primary device
The device in the active-primary state will own the session. If this option is selected and the device that received the first packet is not in the active-primary state, the packet will be forwarded to the peer device over the HA3 link to the owner of the session.

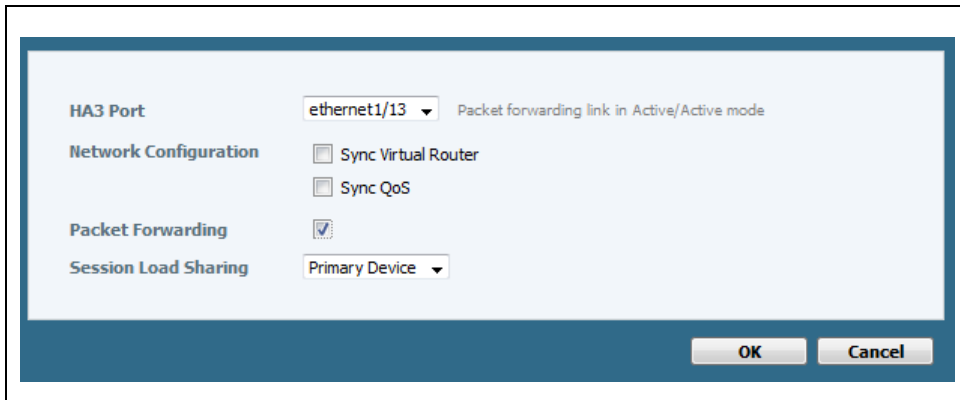
Session setup options

The session setup process is user configurable. One of the following three options can be configured for session setup load sharing

- IP modulo
The session setup load is distributed in the HA cluster by the parity of the source IP address. This option provides a deterministic method of sharing the session setup.
- IP Hash
Hash of either source or combination source/destination IP address is used for distributing session setup
- Primary device
Device in the active-primary state will always setup the session. This setting will result in only one device performing all session setup activities.

Note: Because of the overhead associated with the encapsulation on the HA3 link, the switch ports connecting the HA3 link must be configured to support jumbo frames

If the session ownership is configured to the primary device, then the session setup defaults to the primary device. A sample screen shot of the configuration is shown below.




Note: Setting the session owner and the setup device to be the primary device is not recommended since all traffic processing will be performed by the active-primary device. However for the sake troubleshooting and capturing logs and pcaps it can be useful to set the session owner and setup to primary since the packet processing will not be split between the devices in the cluster. For all other purposes the recommended setting is to use “First Packet” for session owner and “IP modulo” for session setup

Configuring virtual addresses

Case 1: Using floating IP and ARP load sharing

In this example, we configure an A/A cluster on a pair of Palo Alto Networks firewalls, in layer3 mode, using a shared IP for ARP load sharing on the LAN side, and floating IP address on the interfaces connected to the upstream routers.




The table below summarizes the IP addressing scheme

IP address	Interface	Address type	Device binding	Comment
172.35.2.100	ethernet1/1	ARP load sharing	N/A	LAN side- default gateway for hosts
10.1.1.100	ethernet1/2	Floating IP	device-id 0	WAN side
10.1.1.101	ethernet1/2	Floating IP	device-id 1	WAN side

Configuring shared IP- ARP load sharing

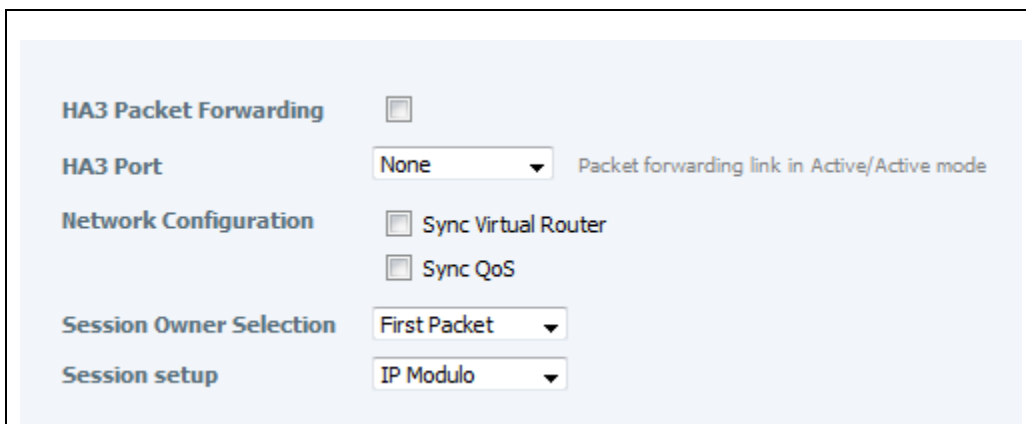
From the **Device>high availability>Virtual address** section click on add to add a new virtual address. From the interface drop down list choose the appropriate interface, and click "Add". From the resulting screen choose the load sharing type to be "arp-load-sharing". In this example we choose ip-modulo as the Load Sharing type.



It is important to note that the ARP load sharing type drop down, determines which one of the devices will respond to the ARP request for the load sharing IP. If the session ownership (configured on the HA3 link option) is set to first packet then the device that responds to the ARP request will become the session owner.

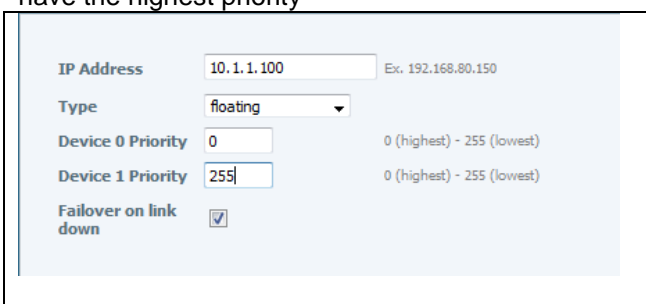
The session setup load sharing will then be determined by the session setup and session load sharing configuration under the HA3 link.

If you select IP-modulo for the ARP load-sharing algorithm as shown above, choose First Packet for session owner, and IP-modulo for the session setup algorithm, as shown in screen shot below, The sessions will be setup and owned by the same device.



Configuring Floating IP

From the **Device>high availability>Virtual address** section click “Add” to add a new virtual address. From the interface drop down list choose the appropriate interface; and click “Add”. From the resulting screen chose the type to be “floating”. The device priority determines which one of the devices will own the floating IP address. We configured two floating IP address, one for each device, with different priorities as shown below. The address with the lower numeric value, will have the highest priority



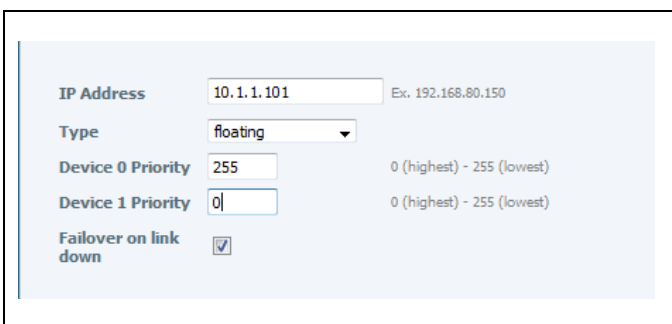
IP Address: 10.1.1.100 (Ex. 192.168.80.150)

Type: floating

Device 0 Priority: 0 (0 (highest) - 255 (lowest))

Device 1 Priority: 255 (0 (highest) - 255 (lowest))

Failover on link down:



IP Address: 10.1.1.101 (Ex. 192.168.80.150)

Type: floating

Device 0 Priority: 255 (0 (highest) - 255 (lowest))

Device 1 Priority: 0 (0 (highest) - 255 (lowest))


Failover on link down:

The command “**show high-availability virtual-address**” can be used to view all configured floating IP addresses

Case2: Using floating IP

In this example we configure floating IP addresses on all of the connected interfaces. Load balancing of the traffic can be accomplished by either using external load balancers or by configuring default gateways on the hosts to each one of the floating IPs. In this example the default gateways on hosts are configured to use the floating IP address

IP address	Interface	Address type	Device binding	Comment
172.35.2.100	ethernet1/1	Floating IP	device-id 0	LAN side
172.35.2.101	ethernet1/1	Floating IP	device-id 1	LAN side
10.1.1.100	ethernet1/2	Floating IP	device-id 0	WAN side
10.1.1.101	ethernet1/2	Floating IP	device-id 1	WAN side



To floating IP address configured can be viewed as shown below

```
admin@PA-1(active-primary)> show high-availability virtual-address
```

```
Total interfaces with virtual address configured: 2
Total virtual addresses configured: 4
```


```
-----
Interface: ethernet1/2                               Virtual MAC: 00:1b:17:00:01:11
  10.1.1.100                                         Active:yes    Type:floating
  10.1.1.101                                         Active:no     Type:floating
-----
Interface: ethernet1/1                               Virtual MAC: 00:1b:17:00:01:10
  172.35.2.100                                       Active:yes    Type:floating
  172.35.2.101                                       Active:no     Type:floating
-----
```

Configuring NAT with active-active cluster

With active-active HA, NAT rules are also matched against the device-id in addition to the IP address and zone information. In the following section we discuss different types of NAT implementation with active-active cluster

Source NAT using interface IP

In this case the source IP address and port numbers of the all connections are translated to the floating IP address configured on the egress interface. The floating IP addresses used for translation are 10.1.1.100 and 10.1.1.101. The hosts are configured with default gateways pointing to each one of the floating IP address to load balance the traffic. It is important to note that two source NAT rules are required, one for each device. However all NAT rules are configured on a single device, which is later synchronized to the peer device



Before configuring the NAT rules, you must verify the floating IP address priorities on the devices. This can be viewed from the Web interface or using the CLI command `>show high-availability virtual-address`

In this example the following priorities are assigned

Virtual Address		
Interface	IPv4	
	Address	Floating
ethernet1/2	10.1.1.100	Device 0 Priority: 0 Device 1 Priority: 255 Failover on link down: ✓
	10.1.1.101	Device 0 Priority: 255 Device 1 Priority: 0 Failover on link down: ✓
ethernet1/1	172.35.2.100	Device 0 Priority: 0 Device 1 Priority: 255 Failover on link down: ✓
	172.35.2.101	Device 0 Priority: 255 Device 1 Priority: 0 Failover on link down: ✓

Note that the floating IP 10.1.1.100 is owned by device-id 0 because of the lowest priority and the IP 10.1.1.101 is owned by device-id 1. The corresponding NAT rules in this case will be

Name	Original Packet						Translated Packet		Active/Active HA Binding
	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation	
src-nat	trust-13	untrust-13	any	any	any	any	dynamic-ip-and-port ethernet1/2 10.1.1.100	none	0
src-nat-1	trust-13	untrust-13	any	any	any	any	dynamic-ip-and-port ethernet1/2 10.1.1.101	none	1

Verification

Let us examine a session flow from the host 172.35.2.4 to a host 10.1.1.250. From the session table, we see that the host 172.35.2.4 is translated to IP 10.1.1.101, the floating IP on PA-2 which is device-id 1

```
admin@PA-4020-100(active-primary)> show session all
-----
ID      Application  State  Type  Flag  Src[Sport]/Zone/Proto (translated IP[Port])
Vsys
-----
19485  telnet      ACTIVE FLOW  NS    172.35.2.4[56484]/trust-13/6 (10.1.1.101[57558])
vsys1
-----
19485  telnet      ACTIVE FLOW  NS    10.1.1.250[23]/untrust-13 (10.1.1.250[23])
vsys1
```

From looking at the session details, we see that PA-2 is the session owner. This is because the first packet of the session from the host 172.35.2.4 was sent to the IP 172.35.23.101 owned by PA-2.

```
admin@PA-2(active-primary)> show session id 19485 | match HA
session synced from HA peer      : False
session owned by local HA A/A    : True
```

The session setup is determined by the session setup load sharing configured on the HA3 link. Note that session setup and the owner device can be the same. In this example the session setup is done by the peer device.

```
admin@PA-2(active-primary)> show counter global filter aspect aa delta yes
Global counters:
Elapsed time since last sampling: 24.406 seconds
name                               value      rate severity  category  aspect
description
-----
ha_aa_session_setup_peer           1          0 info       ha        aa
Active/Active: setup session on peer device
ha_aa_pktfwd_rcv                   1          0 info       ha        aa
Active/Active: packets received from peer device
ha_aa_pktfwd_xmt                   1          0 info       ha        aa
Active/Active: packets forwarded to peer device
-----
Total counters shown: 3
-----
```

From this example we see the session setup device matches the NAT rule bound to the device-id 1, which is the session owner, The source IP translates to floating IP address 10.1.1.101 owned by device-id 1

Source NAT using IP pools

When using address pools with active-active, each device will require its own pool, which will then be associated with a device-id. In this example the following two address pools are created

Address pool name	IP address range	Device binding
Pool-dev0	10.1.1.140-150	Device-id 0
Pool-dev1	10.1.1.160-10.1.1.170	Device-id 1

The corresponding NAT rules will be as shown

Name	Original Packet						Translated Packet		Active/Active HA Binding
	Source Zone	Destination Zone	Destination Interface	Source Address	Destination Address	Service	Source Translation	Destination Translation	
src-nat	trust-13	untrust-13	any	any	any	any	dynamic-ip-and-port Dyn-IP Pool-dev0	none	0
src-nat-1	trust-13	untrust-13	any	any	any	any	dynamic-ip-and-port Dyn-IP-Pool-dev 1	none	1

Note: If you have an existing active-passive configuration with pool based NAT, you will need two separate IP pools in order to implement active-active HA


Source NAT rules with device failure

When one of the devices in a cluster fails;

- All existing sessions will be translated to the current NAT mapping even if the NAT IP is owned by the failed device.
- All new sessions are translated using the NAT rule matching the device-id of the functional device.

Destination NAT in a layer2 connected network

With destination NAT, PAN-OS devices will respond to the ARP request for the destination NAT address with the ingress interface MAC address. In this example, traffic to the destination IP address 10.1.1.200 is translated to 172.35.2.200




When the firewall cluster receives traffic for the destination 10.1.1.200, both devices can respond to the ARP request which may cause network instability. Ideally in such a scenario, only one device will respond. The device in the active-primary state should be configured to respond to the ARP request. The corresponding NAT rule will be as shown below:

Name	Original Packet					Translated Packet		Active/Active HA Binding
	Source Zone	Destination Zone	Source Address	Destination Address	Service	Source Translation	Destination Translation	
DST-NAT Server	untrust-l3	untrust-l3	any	Public server	any	none	address: private server-3	primary

Destination NAT in layer3 connected network

When the firewall cluster is directly connected to layer3 devices as shown in the topology below, both of the firewalls will have to respond to an ARP request for the destination NAT. Traffic can arrive at the firewall from either of the WAN routers.




A destination NAT rule must be configured for both the devices to respond to the ARP request. The NAT rule configured is shown below

Name	Original Packet					Translated Packet		Active/Active HA Binding
	Source Zone	Destination Zone	Source Address	Destination Address	Service	Source Translation	Destination Translation	
DST-NAT Server	untrust-3	untrust-3	any	Public server	any	none	address: private server-3	both

IPSec with active-active cluster

In an active-active cluster the virtual IP address is used as the tunnel end point. In the example below, we establish an IPSec tunnel using the floating IP address of 10.1.1.100 as one of the tunnel end points. When the device owning the floating IP fails, this IP will failover to the peer device in the cluster, still maintaining the tunnel.



Configuration

In the IKE gateway section of the configuration choose floating IP address as the IP address of the interface

IKE Gateway

Local IP Address
 IP
 Floating IP

Peer IP Address Dynamic
Select 'Dynamic' or enter a Peer IP Address

Pre-shared Key

Confirm Pre-shared Key

[Show advanced Phase1 options...](#)

A tunnel interface must be configured on each of the devices, with a unique IP address. Note that the IP addresses of the interfaces are not synchronized between the peers. The routes via the tunnel interface will be synchronized between devices.

Tunnel status

It is important to note that the device in the cluster that owns the floating IP address will have both the Phase1 and Phase2 SA established. The peer device will however only have the phase2 SA. In our example, we see the PA-2 owns the floating IP address 10.1.1.100 used as the tunnel end point.




```
admin@PA-2 (active-secondary)> show high-availability virtual-address
```

```
Total interfaces with virtual address configured: 2
Total virtual addresses configured: 4
-----
Interface: ethernet1/2                               Virtual MAC: 00:1b:17:00:01:11
  10.1.1.100                                         Active:yes      Type:floating
  10.1.1.101                                         Active:no       Type:floating
-----
Interface: ethernet1/1                               Virtual MAC: 00:1b:17:00:01:10
  172.35.2.100                                       Active:no       Type:floating
  172.35.2.101                                       Active:yes      Type:floating
-----
```

PA-2: Owns the floating IP

IKE Gateway							Tunnel Interface			
Name	Status	Interface	Local IP	Peer IP	Status	Interface	Virtual Router	Security Zone	Status	
<input type="checkbox"/> VPN-1		ethernet1/2	10.1.1.100 (floating-ip)	10.1.1.250		tunnel.1	default (Show Routes)	trust-I3		

PA-1

IKE Gateway							Tunnel Interface			
Name	Status	Interface	Local IP	Peer IP	Status	Interface	Virtual Router	Security Zone	Status	
<input type="checkbox"/> VPN-1		ethernet1/2	10.1.1.100 (floating-ip)	10.1.1.250		tunnel.1	default (Show Routes)	trust-I3		


Also note the encap and decap byte counter increment only on the firewall PA-2, which owns the floating IP address

```
admin@PA-2(active-secondary)> show vpn flow name VPN-1 | match bytes
encap bytes:          349544
decap bytes:          350608
admin@PA-4020-151(active-secondary)>
```

```
admin@PA-1(active-primary)> show vpn flow name VPN-1 | match bytes
encap bytes:          0
decap bytes:          0
```

SSL VPN

The floating IP address can be used as a tunnel termination IP address of SSL VPN connection. The device that owns the floating IP will terminate the SSL VPN connections. The user authentication table is synchronized between both the devices in the cluster.



In this example, PA-2 firewall owns the floating IP address.

Verifying SSL flow

The command “`show ssl-vpn flow name` ” is used to verify the SSL VPN flow.

```
admin@PA-2(active-secondary)> show ssl-vpn flow name SSL_VPN_Cluster
tunnel  SSL_VPN_Cluster
      id:                5
      type:              SSL-VPN
      local ip:         10.1.1.100
      inner interface:  tunnel.2           outer interface: ethernet1/2
      ssl cert:        SSL-VPN
      active users:    1
assigned-ip      remote-ip      encapsulation
-----
172.16.2.1      10.1.1.98      IPSec SPI 3AD402B6 (context 7)
```

The remote-ip column shows 0.0.0.0 if the device does not own the floating IP address

```
admin@PA-1(active-primary)> show ssl-vpn flow name SSL_VPN_Cluster
tunnel  SSL_VPN_Cluster
      id:                5
      type:              SSL-VPN
      local ip:         10.1.1.100
      inner interface:  tunnel.2           outer interface: ethernet1/2
      ssl cert:        SSL-VPN
      active users:    1
assigned-ip      remote-ip      encapsulation
-----
172.16.2.1      0.0.0.0        IPSec SPI 3AD402B6 (context 7)
```

Logs and packet capture

All traffic logs are logged by the session owner. When the session owner fails, the other device will become the session owner and will handle logging. Should the failed device recover before the session ends, it will take back ownership of the session and handle logging.

Summary

The fundamental architecture of Palo Alto Networks’ next-generation firewalls provides network security by enabling enterprises to see and control applications, users, and content – not just ports, IP addresses, and packets – using three unique identification technologies: App-ID, User-ID, and Content-ID. These features are complimented by robust hardware which is designed to provide separation of control plane and data plane This separation ensures highly available management, high speed logging and route updates, while the dataplane does all the security processing. The Active-Active HA cluster supports implementing the Palo Alto network firewall in high performance, scalable networks.