How to Configure and Use High Availability

https://campus.barracuda.com/doc/26576260/

For redundancy and reliability, you can set up two Barracuda NextGen X-Series Firewalls in a high availability (HA) cluster. During normal operations, the primary unit is active while the secondary unit waits in standby mode. The secondary unit has the same configurations as the primary unit, and it only becomes available when the primary unit is down. The failover is reversed when the primary unit can resume operations. Services should be configured on the secondary IP address, not the management IP address of the firewall, because only the secondary IP addresses fail over to the secondary unit. For the same reason, use the secondary IP address as the default gateway for your clients.

To execute a failover when a unit or networking component becomes unavailable, you can configure the monitoring of additional IP addresses and interfaces. You can also manually execute a failover.

When installing two firewalls in a high availability cluster, ensure that the cabling is done exactly the same on both units. The management IP addresses must also be configured on the same ports. For example, if port 3 on the primary box is connected to ISP 1, the secondary box must also connect port 3 with ISP 1. If you install cabling incorrectly, HA failover does not work properly. For an example of correct cabling, see the following diagram:

Before You Begin

- If you want to join a Windows domain, you must do so on both primary and secondary units before creating the HA cluster. For more information, see How to Join a Windows Domain.
- If you want to use the Barracuda Web Security Service, you must connect both primary and secondary units before creating the HA cluster. For more information, see Cloud Features.
- Each X-Series Firewall must have a management IP address in the same subnet. Verify that they are not using the same IP addresses as the management IP address.

Step 1. Add Management IP Addresses to the Administrator IP/Ranges

If you restrict administrative access to the firewall by defining administrator IP addresses or networks, you must add the management IP address of the HA partner unit to the administrator IP/Ranges list. If you are not restricting the administrator IP address (0.0.0.0 entry is present), you can skip this step.

Step 1.1 Administrator IP/Range on the Primary Unit

Add the management IP of the secondary unit to the administrator IP addresses on the primary unit.

1. Log into the primary unit.
2. Go to **BASIC > Administration.**
3. In the **ADMINISTRATOR IP/RANGE** section, enter:
    - **IP/NETWORK ADDRESS** – Enter the management IP address of the secondary unit.
    - **NETMASK** – Enter 255.255.255.255
4. Click **ADD**.

Step 1.2 Administrator IP/Range on the Secondary Unit

Add the management IP of the primary unit to the administrator IP addresses on the secondary unit.

1. Log into the secondary unit.
2. Go to **BASIC > Administration.**
3. In the **ADMINISTRATOR IP/RANGE** section, enter:
    - **IP/NETWORK ADDRESS** – Enter the management IP address of the primary unit.
    - **NETMASK** – Enter 255.255.255.255
4. Click **ADD**.

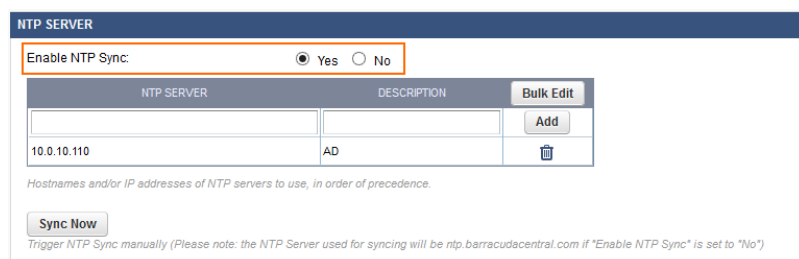## Step 2. Add a Secondary IP Address to the Primary Firewall

Add a secondary IP address to the primary firewall and configure the services of the firewall that are to be used from the local network to listen on this IP address. Use this secondary IP address as the default gateway for the clients in your network. In case of a failover, this IP address is transferred to the secondary firewall.

1. Go to **BASIC > IP Configuration**.
2. Enter a **Secondary IP Address** and select the services that should listen on this IP address.
3. Click **Add**.

Step 3. Enable NTP

Go to **BASIC > Administration** and verify that NTP is enabled on the primary unit.

Step 4. Enable High Availability

Before you set up two X-Series Firewalls in an HA cluster, ensure that both units fulfill the following prerequisites:

- Both firewalls must be the same model type and revision. They must also run the same firmware version.
- The management IP addresses of both units must be in the same network and subnet.
- System clocks and time zones must be accurately set on both units. If they are not, HA pairing can fail.
- The **Default Domain** (**BASIC > Administration**) must be set on both units.

Enable HA on the Secondary Unit

1. Log into the secondary unit.
2. Go to **ADVANCED > High Availability**.
3. In the **Setup** section, click **Enable High Availability**.
4. In the **Enable High Availability** window, enter the management IP address, serial number, and administrator password for the primary unit.
5. Click **Enable**. The HA pairing process can take several minutes. During this process, do not

reload the configuration page or configure any other settings.

After the HA pairing is successful, the **Disable High Availability** option appears in place of the **Enable High Availability** option. The IP addresses and serial numbers of both HA units are also displayed.

Additionally, this warning message is displayed on every configuration page of the secondary unit:

Warning: Attention! This is the secondary High Availability box. Go to the primary box to edit the configuration.

While the secondary unit is part of the HA cluster, you can configure only the following settings:

- **ADVANCED > High Availability**
- **NETWORK > IP Configuration > Management IP Configuration**
- **NETWORK > IP Configuration > Dynamic Interface Configuration**
- **(If 3G is available) NETWORK > IP Configuration > 3G Network Interface**

Configure Monitoring

You can configure the monitoring of additional IP addresses and interfaces. If these IP addresses and interfaces become unreachable, a failover is executed.

On the **ADVANCED > High Availability** page, in the **Monitoring** section, add the **Reachable IPs** and **Reachable Interfaces**.

Verify the HA Status

To verify the HA status of the firewall, go to the **ADVANCED > High Availability** page and see the **Status** section. This section indicates if the appliance is active, standby, primary, or secondary. If the appliance is not part of an HA cluster, this section indicates that it is **Stand-Alone**.

This figure shows an example of the status for a firewall in a high availability cluster.

On the **BASIC > Status** page, you can also view the current HA status in the **Services** section. To see the status details, hover over **High Availability**.

Note that the secondary X-Series Firewall is not visible in Barracuda Cloud Control.

Manually Execute an HA Failover

On the **ADVANCED > High Availability** page, you can manually execute an HA failover by clicking **Manual Failover** in the **Status** section of the unit that is currently active.

If the X-Series Firewall is not part of an HA cluster, the **Manual Failover** option is disabled.

Settings Not Synced Between Units in a High Availability Cluster

The following settings are unique to each unit in the high availability cluster and are not synced:

- Domain
- Hostname
- Timezone
- HTTPS Port
- Management Interface Configuration.
- Content of DNS Cache
- Dynamic Interfaces

Figures

1. ha_conf_01-01.png
2. HA_NTP.png
3. image2013-7-26 15:21:15.png
4. HA_status01.png