High Availability

https://campus.barracuda.com/doc/41115943/

A standalone system is typically set up in an HA cluster from one of the following configuration scenarios:

- It is an existing standalone Barracuda NG Firewall, to which a second NG Firewall is added for high availability.
- It is one of two existing standalone Barracuda NG Firewalls that are to be configured into a single HA pair.
- It is part of an HA pair that is to be installed from scratch. In this case, install the new system and then set it up in HA mode.

It is important to configure switches and routers properly to work in conjunction with an HA setup. Most important is the ARP cache time or ARP timeout of the networking equipment. When the secondary unit starts its services, it uses the same IP addresses (except for the management IP address) as the primary unit, but with different MAC addresses. With an infinite timeout configured, the secondary unit would never be reached, because the MAC address would be resolved to the wrong port. With a timeout of 300 seconds, the secondary unit would not be reached for 5 minutes and the HA concept would not fulfill its purpose. The recommended setting lies between 30 and 60 seconds. Also note that the number of ARP requests will increase with a higher timeout.

In this article:

HA Monitoring without a Private Uplink

In an HA system with no private uplink, alive packets and status information are transferred over the network that the management IP addresses belong to. For example, in the following diagram, the HA state is exchanged via the 10.0.8.0/24 network.

When the switch "dies", the connection between the HA partners also breaks, and the secondary unit starts its servers although the primary unit is still alive. When the switch is reactivated, for around 1 second, both units are up and duplicate IP addresses are online until the primary unit stops its servers.

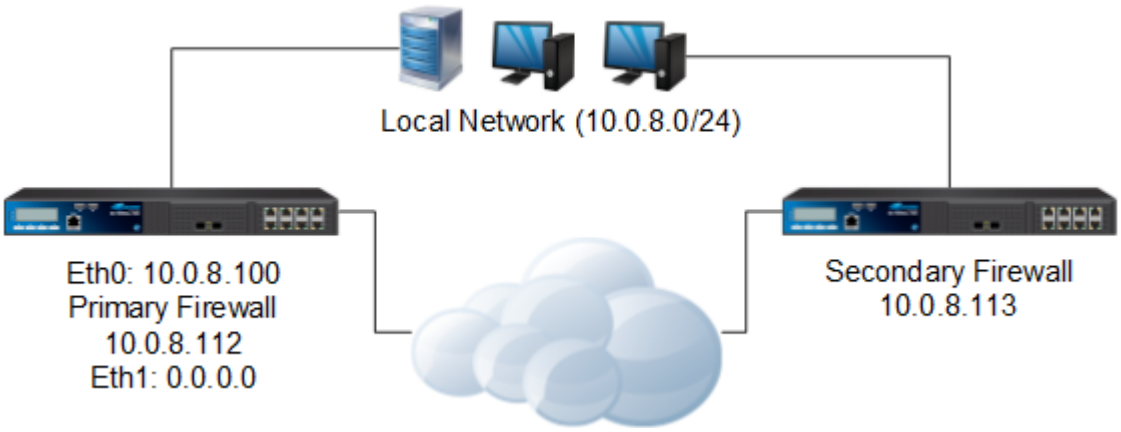

HA Monitoring with a Private Uplink

In an HA cluster with a private uplink, one network interface is dedicated for HA purposes. An example of this setup is displayed in the figure below. There are some routing specialties (host routes) to route the HA traffic via the private uplink. A failover route must also be configured to make sure that the units can reach each other via both routes. The private uplink should be a direct connection with a crossover cable so that it is independent from a further hardware component (switch/HUB). The subnet for the uplink should be a 2-bit network.

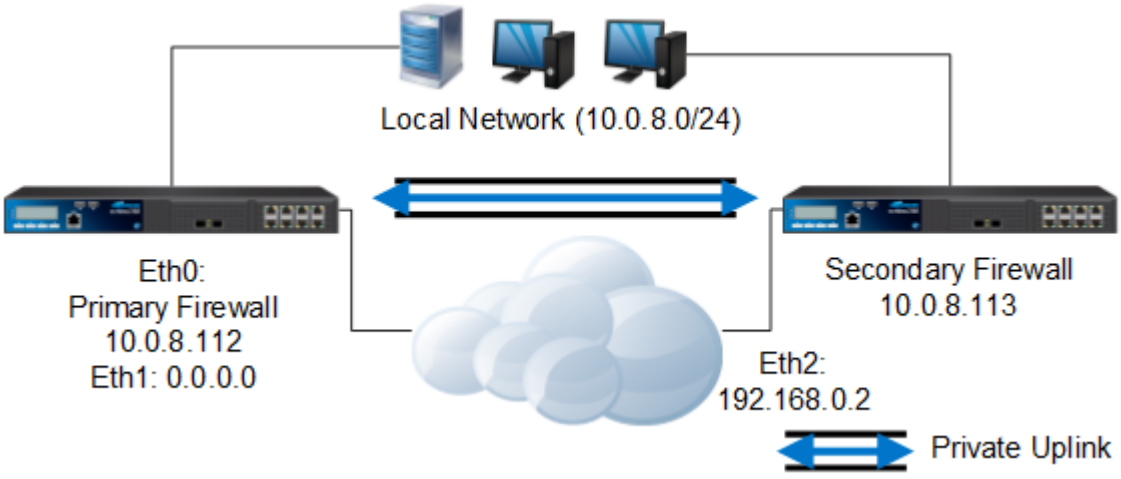

Designing an HA System

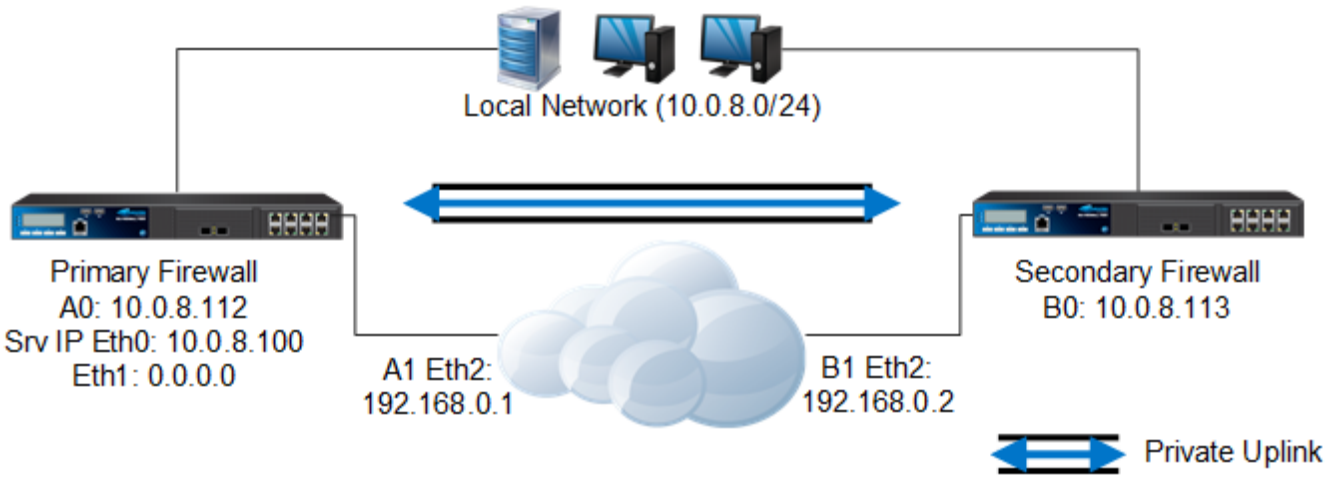

Example IP Addresses

	Primary Unit	Secondary Unit
Management IP	10.0.8.112 / eth0	10.0.8.113 / eth0
FW Server IP	10.0.8.100	
Further Network (Private Uplink)	192.168.0.1/30 / eth2	192.168.0.2/30 / eth2

The route the heartbeat takes is configured via the parameter group **Translated HA IP** (**Config** > **Box** > **Infrastructure Services** > **Control**). In the example settings, the heartbeat is configured to use both the 10.0.8.0/24 network AND the private uplink to send heartbeats.

	Translated HA IP	Alternative HA IP	Usage Policy
Primary Unit	10.0.8.113	192.168.0.2	Use-Both
Secondary Unit	10.0.8.112	192.168.0.1	Use-Both

Configure the Translated HA IP and Alternative HA IP on the primary and secondary unit. These IP addresses are used in the default firewall rules for HA synchronization that allow HA traffic between both HA partners.

The HA IP address must be a Management IP address. Otherwise, the control daemon does not listen on the alternative HA IP, causing heartbeat and sync to fail.

If you are running an HA setup with different appliance revisions, ensure that both physical ports of the private uplink are using identical port labels. Otherwise, HA synchronization may fail.

Figures

1. ha_sys.0.png
2. ha_sys.png
3. ha_sys1.png