Planning Your High Availability Deployment

https://campus.barracuda.com/doc/5799967/

Clustered Barracuda Link Balancers may require extra equipment.You may need to add switches so that the WAN links can connect to two systems. To deploy in front of an existing firewall, you will need to add a switch between the Barracuda Link Balancers and the firewall (or two switches for dual firewalls).

The following figures show sample deployments of a pair of clustered Barracuda Link Balancers with: a single firewall; two clustered firewalls; and with no external firewall.

In Front of a Single Network Firewall

The figure below shows two Barracuda Link Balancers deployed with one network firewall. The LAN IP addresses of the two Barracuda Link Balancers and the VRRP virtual IP address must all be on the same subnet.
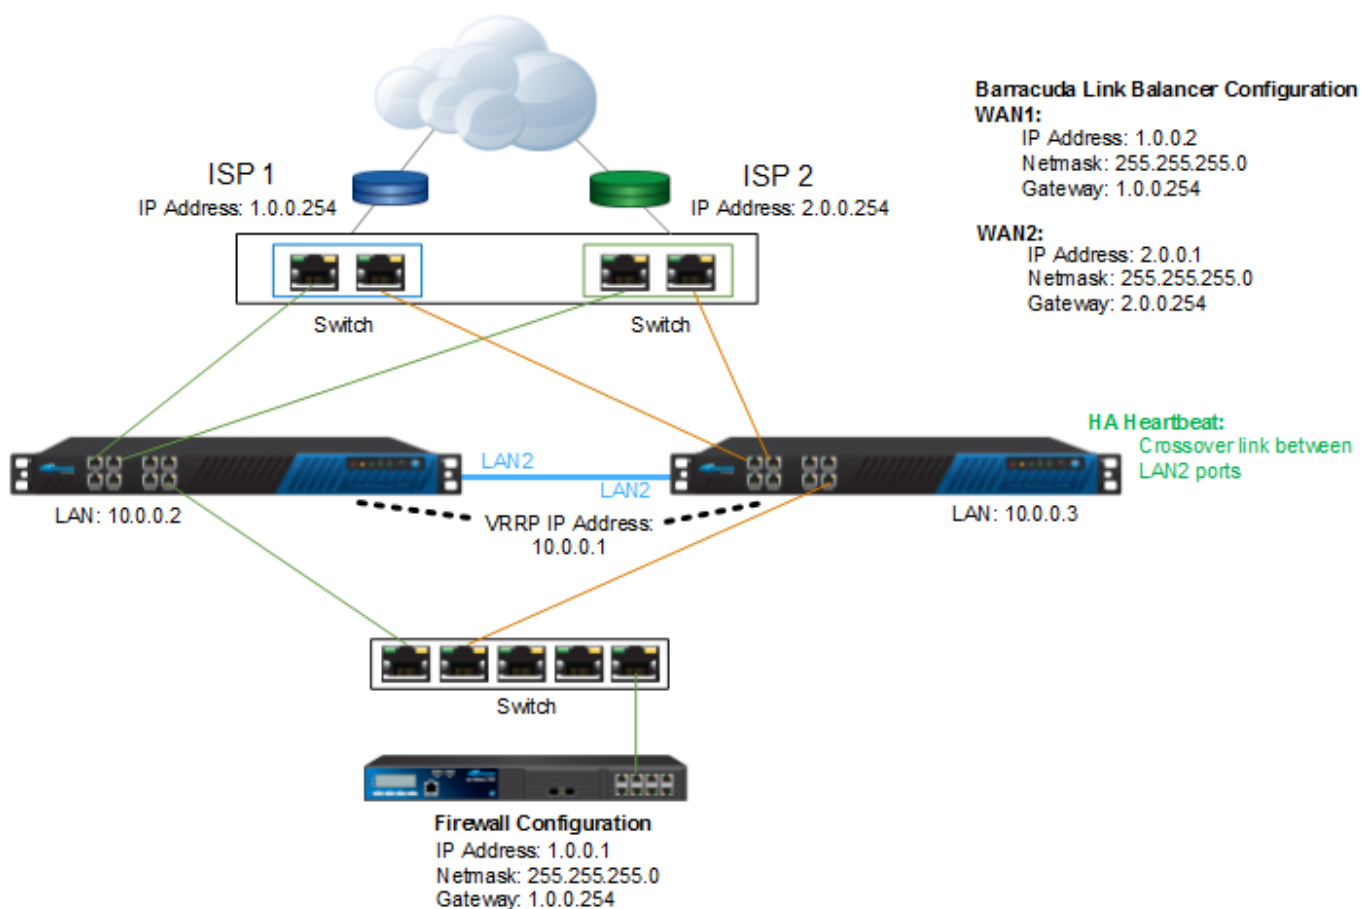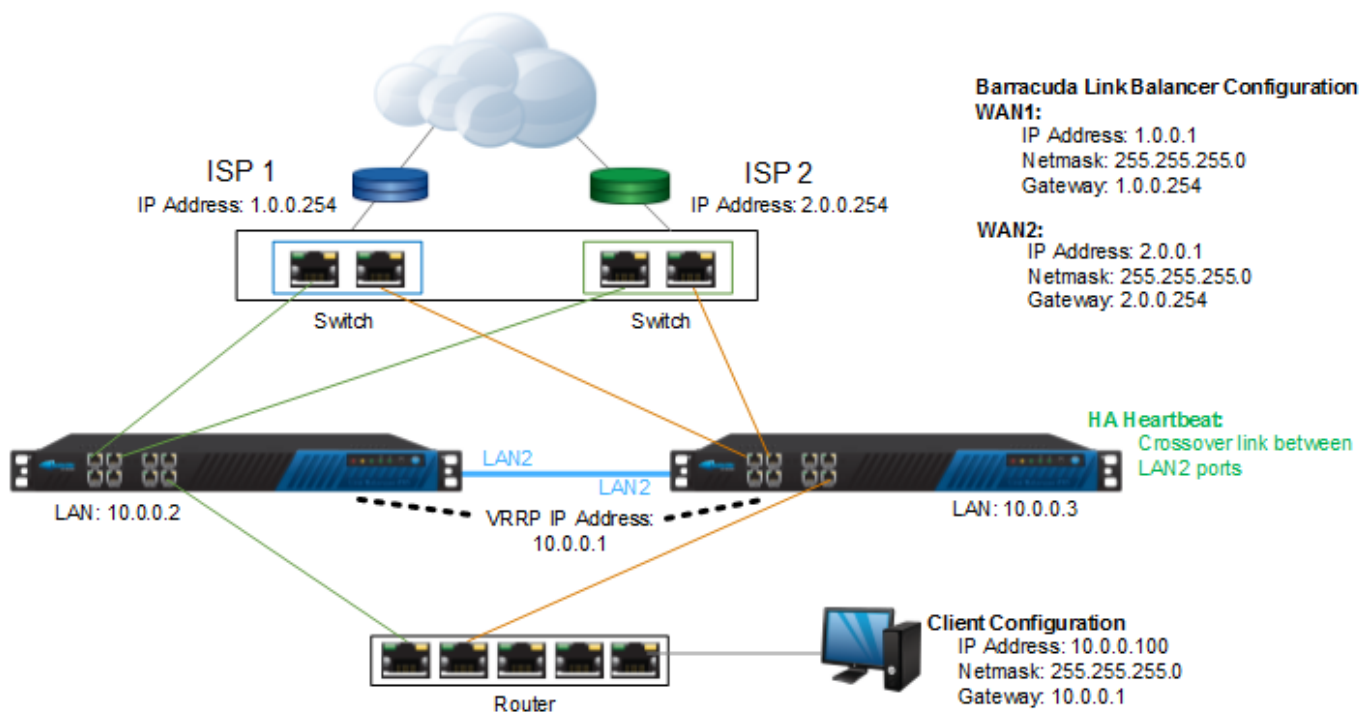
In Front of Dual Network Firewalls

The following figure shows two Barracuda Link Balancers and two clustered firewalls. The LAN IP addresses of the two Barracuda Link Balancers and the VRRP virtual IP address must all be on the same subnet.
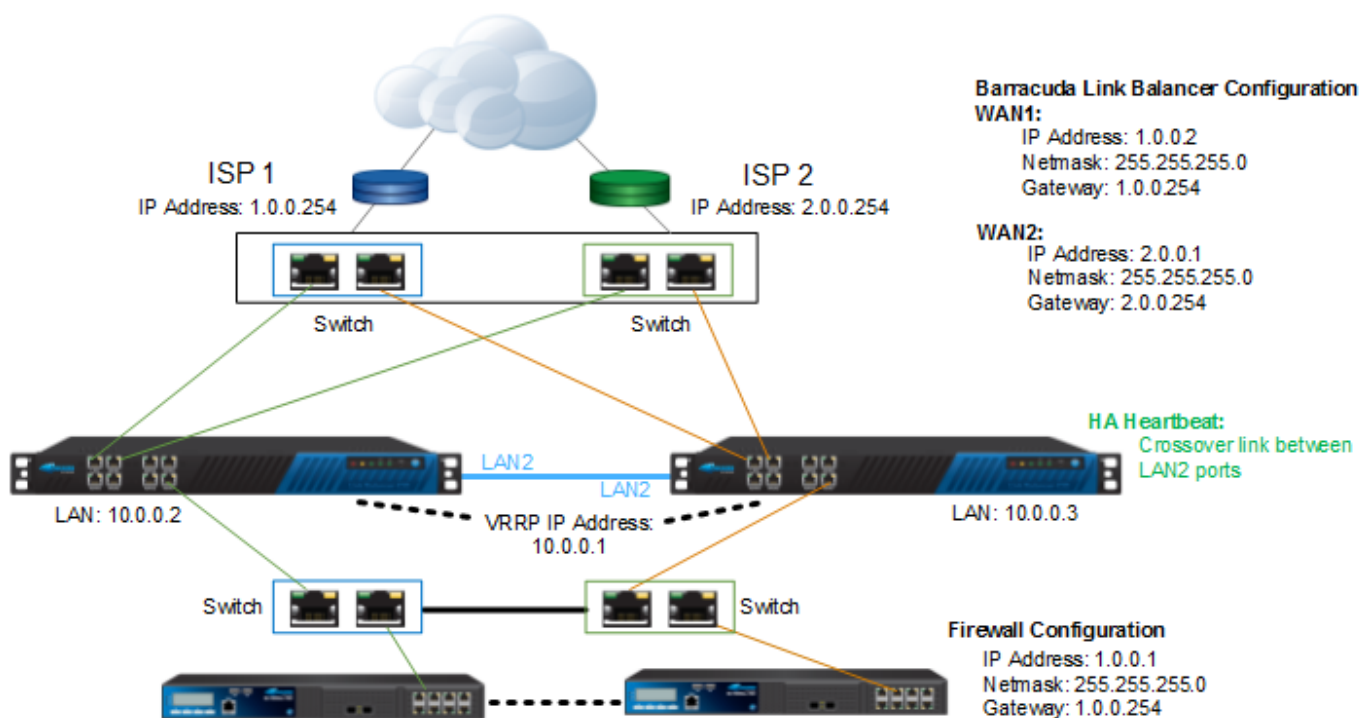
No External Firewalls

The following figure shows two Barracuda Link Balancers with the firewall enabled. As in the other deployment examples, the LAN IP addresses of the two Barracuda Link Balancers and the VRRP virtual IP address must all be on the same subnet. Note that only in this example, the VRRP virtual IP address is the default gateway for devices on the LAN. If you add a second Barracuda Link Balancer to a network where the gateway of the client devices was already configured to use the LAN IP address of the first Barracuda Link Balancer, you could assign a new LAN IP address to that Barracuda Link Balancer and use its original LAN IP address as the VRRP virtual IP address.

Barracuda Link Balancer Configuration
WAN1:
IP Address: 1.0.0.1
Netmask: 255.255.255.0
Gateway: 1.0.0.254

WAN2:
IP Address: 2.0.0.1
Netmask: 255.255.255.0
Gateway: 2.0.0.254

HA Heartbeat:
Crossover link between LAN2 ports

Client Configuration
IP Address: 10.0.0.100
Netmask: 255.255.255.0
Gateway: 10.0.0.1

Best Practices for Setting Up HA on Models 330 and 430

- Use switches (resulting in multiple collision domains) instead of hubs (resulting in a single collision domain).
- Make sure the ARP cache timeout on the switches is not set too high. Typically, 60 to 180 seconds is a good range.
- You should use different switches for the various WAN links instead of one common switch for all of them. However, if only a single common switch is used, make sure to use a manageable switch and configure your VLANs to create multiple broadcast domains for each WAN link respectively.
- Make sure your network architecture logically matches the respective figure on this page.

Figures

1. in_front.png
2. in_front_dual.png
3. ha_no_fw.png