How to Manage a High Availability Environment with Two Barracuda Load Balancers

https://campus.barracuda.com/doc/3538993/

For an overview of High Availability, and a list of requirements, see the article Understanding Barracuda Load Balancer High Availability.

Manage Access to the Passive System

Unless you wish to use the WAN IP address to access the web interface of the passive system, configure a management IP address in the **Management IP Configuration** section on the **BASIC > IP Configuration** page. This management IP address is used by the Ethernet port on the back of the passive system.

Failover if LAN Link Goes Down

There is an option to fail over to the passive system if the active system cannot detect its LAN link. In one-armed deployments (including Direct Server Return), the LAN port does not need to be monitored as the Real Servers are all connected to the WAN, so this option *should be disabled*. If the Barracuda Load Balancer is in bridge-path mode, LAN port monitoring is compulsory.

Forceful or Manual Failover

You can force failover to the passive system using the web interface. This transfers the load to the passive system without bringing down any of the interfaces of the active system. When the passive system has become active, LAN or WAN cables can be removed or other maintenance performed on the now-passive system.

Primary and Backup Roles

When two systems are joined in a cluster, the system that joins the cluster is the backup system. The other one has the role of primary system. Initially, the primary system is the active system. Either of the systems in a cluster is capable of being the active system. The backup and primary roles are important when discussing failback.

Failback

There is an automatic failback option that can be configured if you want the originally active (primary) system to take over the Virtual IP addresses and resume load balancing upon its recovery after a failover. This option can be found on the **ADVANCED > High Availability** page.

You can manually switch to the primary system using the Failback command that is available on the same page.

It may be better to opt for manual failback, as it can minimize the number of times that service is interrupted. For example, if the primary system suffers an outage, the backup system takes over. When the primary system recovers, if automatic failback is selected, then it will once again become the active system. This means two interruptions of service. If manual failback is selected, then the backup system will continue processing traffic even after the recovery of the primary system.

Synchronize Data between Clustered Systems

When two Barracuda Load Balancers are initially joined, most configuration settings are copied from the primary system in the cluster to the backup system (the system that joins the cluster). These settings are synchronized between the systems on an ongoing basis.

The following data is *shared* between the clustered systems:

- Global system settings configured through the web interface
- Any installed SSL Certificates
- All static routes and VLANs, etc., configured on the **ADVANCED > ADVANCED IP Config** page

The following data is *unique* between the clustered systems:

- All of the system IP configuration (WAN IP address, management IP address, operating mode, DNS servers and domain) configured on the **BASIC > IP Configuration** page except for the LAN IP address
- System password, time zone, and web interface HTTP port as configured on the **BASIC > Administration** page
- Parameters on the **ADVANCED > Appearance** page
- The HTTPS port and SSL certificate used to access the web interface as configured on the **ADVANCED > Secure Administration** page

Specify Source IP Address in a Clustered Environment

By default, the source IP address of traffic sent to the Real Servers is translated (source NAT'd) to be the WAN IP address of the Barracuda Load Balancer. If the Barracuda Load Balancer is clustered, the WAN IP address is not shared between the two clustered systems. To use the same source IP address in the event of failover, complete the following steps on the active system to create a custom virtual interface that associates an externally-accessible IP address with the WAN port. Use this IP address to create a source NAT rule. This interface is used by the backup system if failover occurs. The changes made are propagated automatically to the passive system.

1. On the **ADVANCED > Advanced IP Config** page, in the **Custom Virtual Interfaces** table, create a custom interface by filling in the following fields:
    - **IP/Network Address** and **Netmask** – an externally-accessible IP address and netmask (may be on the same subnet as the WAN IP address)
    - **Service/Interface Name** – e.g., SNAT IP Address
    - **Port** – Select WAN from the list
2. In the **Source Network Address Translation** table, create a source NAT rule by entering values for:
    - **Internal Address**, **Netmask** – Enter an internal IP address and netmask. Usually this is the Real Server network
    - **Outbound Address** – The externally-accessible IP address from step 1.
    - **Port** – Select the WAN port

Update Firmware on Clustered Systems

Update the passive system first. On the passive system:

1. If the passive system is also the backup system, go to the **ADVANCED > High Availability** page and set **Failback Mode** to **Manual**. This ensures that the transfer of control from one system to the other does not happen while the firmware update is in process.
2. Go to the **ADVANCED > Firmware Update** page and follow the instructions there to update the firmware.

When the update is complete and the passive system is running the new firmware, update the firmware on the active system.

When all updates are complete, you can failback to the primary system, if desired, by either:

- Changing **Failback Mode** to **Automatic**, or
- Clicking **Failback** in the Clustered Systems table

Related Articles

- Understanding Barracuda Load Balancer High Availability
- How to Configure the Barracuda Load Balancers for High Availability
- How to Remove a Barracuda Load Balancer from a High Availability Environment