

Understanding Barracuda Load Balancer High Availability

<https://campus.barracuda.com/doc/3539028/>

High availability (HA) is an advanced feature and may not be appropriate for all environments; contact [Barracuda Networks Technical Support](#) *before enabling this feature*.

The High Availability (HA) environment is available on the Barracuda Load Balancer 340 and higher.

- If the Server is in the same network as the custom virtual interface, then the custom virtual interface is used to connect to the Server using the interface route/static route or the default gateway, in that order.
- If the Server, the custom virtual interface, and the WAN IP are all in the same network, you cannot use the custom virtual interface to connect to the Server. In this scenario, the WAN IP is always used to connect to the Server.
- The virtual interface of the service can be in any network.

High Availability Operation

Use High Availability (HA) to cluster two Barracuda Load Balancers as an active-passive pair. Only one system actively processes traffic at any one time, but the two systems continuously share almost all configuration and monitor each other's health.

The active system in a clustered pair handles all of the traffic until one of the following conditions is encountered:

- Passive system detects that the active system is no longer responsive on the WAN;
- Active system detects that its LAN connection has been lost;
- Administrator manually forces failover using the web interface;
- Active system encounters a hardware failure (including a power failure) or a failure in one of its critical software modules.

If any of these conditions is encountered, the passive Barracuda Load Balancer then:

- becomes active;
- assumes all of the Virtual IP addresses of the services and the LAN IP address of the other Barracuda Load Balancer;
- performs the load balancing.

Clustered Barracuda Load Balancers negotiate which is the active one according to the Virtual Router Redundancy Protocol (VRRP) specification. The two systems must be configured with the *same cluster shared secret and group ID*. If other systems on the same subnet are also using VRRP, the cluster group ID *must be unique*.

The passive Barracuda Load Balancer does not do any load-balancing or monitoring of Services or Real Servers. For example, in the web interface of the passive system, all of the Services and Real Servers on the **BASIC > Services** have red health indicators.

Requirements

Before joining two systems, each Barracuda Load Balancer must meet the following requirements:

- Barracuda Load Balancer 340 or higher; note that *both Barracuda Load Balancers must be the same model*;
- Activated and on the same version of firmware;
- Able to access all Real Servers;
- On the same physical network segment;
- Able to reach the other Barracuda Load Balancer on the WAN interface.

When clustering two Barracuda Load Balancers for high availability, Barracuda Networks recommends enabling PortFast on the relevant ports.

In addition, the active system should be fully configured; see [Services](#) and [Configuring the Load Balancer Network](#) for a complete list of service and network configuration tasks. For the passive system, follow the instructions in the [Configuring the Passive System](#) section. Once the systems are clustered, you must manually reboot the passive system.

Do not configure Services on the passive system.

To speed up recognition of a newly active Barracuda Load Balancer, disable spanning tree protocol on the ports of the switch where the WAN ports of the two Barracuda Load Balancers are connected. If it is a Cisco switch, enable Spanning Tree PortFast on the ports connected to the WAN ports of the Barracuda Load Balancers.

When the Barracuda Load Balancer becomes active it sends out a gratuitous address resolution protocol (ARP). It continues to send a gratuitous ARP every minute; the passive system *does not* issue any ARPs.

Related Articles

- [How to Configure the Barracuda Load Balancers for High Availability](#)
- [How to Manage of a High Availability Environment with Two Barracuda Load Balancers](#)
- [How to Remove a Barracuda Load Balancer from a High Availability Environment](#)

© Barracuda Networks Inc., 2024 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.