

Virtual Server Monitoring

<https://campus.barracuda.com/doc/73718970/>

To ensure and maintain the connectivity of a virtual server, you can define pools of IP addresses and/or network interfaces that are continuously monitored by the Barracuda CloudGen Firewall. If the health check of a monitored IP address or the link state of a network interface fails, the virtual server is automatically shut down. As soon as the health check target is successful, the virtual server is started again. Monitoring policies define which requirements must be met for the virtual server to remain active, or to be shut down. If you are using an HA cluster, you can use monitoring policies to define the behavior of the secondary HA unit. If necessary, you can use custom scripts which are executed when the virtual server is started or stopped.

Layer 3 Monitoring

The Layer 3 monitoring policy defines the settings for IP address monitoring. The policy configuration provides two address pool tables. Add the target addresses to the tables. These IP addresses must be reachable for the virtual server to stay up. The following Layer 3 monitoring policies are available:

- **all-OR-all-present** – All of the IP addresses from at least one IP address pool, e.g., from the **Monitored IPs I** table, must be reachable. If you enter IP addresses in both the **Monitored IPs I** and **II** tables, the IP addresses from at least one of these tables must be available. Otherwise, the virtual server is deactivated.
- **one-AND-one-present** – At least one IP address from each monitoring pool must be reachable. If you only enter IP addresses in the **Monitored IPs I** table, at least one IP address from this table must be available. If you enter IP addresses in both tables, at least one IP address in each table must be available.

The control service runs an ICMP check on all IP addresses in 10-second intervals. If no answer is received, the IP addresses are probed every second for a 10-second period. If no response is received from a valid health check target during the 10-second period, the virtual server shuts down. The server is reactivated as soon as an answer is received for the subsequent probes. If the monitor target is part of a directly attached subnet, ARP responses are monitored too. If the ICMP response fails with presence of a valid ARP entry, the server is still marked as healthy.

Example Setup:











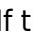

Layer 3 monitoring is configured for the virtual server S2, using both address pools with the following IP addresses and statuses:

Monitored IPs I	Status	Monitored IPs II	Status
10.0.10.110	up	10.0.10.88	up













10.0.10.68	down	10.0.10.99	down
------------	------	------------	------

The status of the virtual server is displayed on [Control > Server](#) page:

If the monitoring policy **one-AND-one-present** is used, the server stays up because one IP address of each address pool is available.

Server	Status	Statu...	IP Addresses
 S2	primary	block	10.0.10.67,172.16.0.50,62.99.0.110
 Status	primary		
 HA	block		
 Monitor S2	OK	OK	
 IP	OK	N.A.	
 list1	OK	N.A.	
 10.0.10.110	OK	N.A.	
 10.0.10.68	FAIL	N.A.	
 list2	OK	N.A.	
 10.0.10.88	OK	N.A.	
 10.0.10.99	FAIL	N.A.	
 IP			

If the **all-OR-all-present** policy is used, the server shuts down because at least no IP pool is fully available.

Server	Status	Statu...	IP Addresses
 S2	disabled	down	10.0.10.67(!),172.16.0.50(!),62.99.0.110
 Status	disabled		
 HA	down		
 Monitor S2	FAIL	OK	
 IP	FAIL	N.A.	
 list1	FAIL	N.A.	
 10.0.10.110	OK	N.A.	
 10.0.10.68	FAIL	N.A.	
 list2	FAIL	N.A.	
 10.0.10.88	OK	N.A.	
 10.0.10.99	FAIL	N.A.	
 IP			

Layer 2 Monitoring

The Layer 2 monitoring policy defines the settings for interface monitoring. Add the interfaces that should be checked according to the policy in the **Monitored Interfaces I** and **II** tables. Layer 2 monitoring is available in the following modes:

- **all-OR-all-present** – All of the interfaces from at least one interface pool, e.g. from the

Monitored Interfaces table, must be available.

- **one-AND-one-present** – At least one interface from each interface pool table must be available. If you have added interfaces in one table, at least one IP address from this table must be available. If you have added interfaces in both tables, at least one interface from each table must be available.

The control service checks the link status of each interface on a regular basis. Depending on the selected policy, the server is shut down if the links on the monitored interfaces are unavailable. The server is restarted when the links of the monitored interfaces are up again.

Server Monitoring in High Availability Clusters

If your Barracuda CloudGen Firewall is part of an HA cluster, you can extend the monitoring policy to both units. For HA monitoring, you can select the following options:

- **Monitoring on Backup Box** – If set to **No** (default), server monitoring on box and HA box is processed only by the primary unit. In case of failover, the non-availability of health check targets is ignored by the HA box and the server stays up on the secondary unit. If set to **Yes**, the monitoring policy will also be enforced by the backup box. In case of a failover, the virtual server is then also deactivated on the second unit if the monitoring also fails on the secondary unit.
- **Shared-HA-Probing** – Shared HA probing combines the IP address and interface information of both units. Both sets of IP addresses or interfaces must be available on both units. An IP address or interface that is not operational on both HA peers will be excluded from the HA logic decision. If a server is active on a unit and blocked on the peer unit, any probing results will be ignored. The probing decision will only be made if a situation persists over two probing cycles. This gives the system time to account for the delay between detection and synchronization and avoids aliasing effects.
- **Local-HA-Probing** – (default) Only local health check target resources are probed. This means every HA partner performs its own monitoring procedure.

Step 1. Configure the Operation Mode

Configure the monitoring policies for IP addresses and interfaces that must be reachable in order for the virtual server to stay up. When your Barracuda CloudGen Firewall unit resides in an HA cluster, specify the monitoring policy for the case of HA failover:

1. Go to **CONFIGURATION > Configuration Tree > Box > Virtual Servers > your virtual server > Server Properties**.
2. In the left menu, select **Monitoring Policy**.
3. Click **Lock**.

4. From the **Monitoring on Backup Box** list, select whether monitoring should be performed and, in case of failover, adapted by a secondary HA unit.
5. Select the **Probing Policy**. For more information, see [Server Monitoring in HA Clusters](#).

Step 2. Configure the Monitoring Policy

Specify the monitoring policy for IP addresses and interfaces.

1. Go to **CONFIGURATION > Configuration Tree > Box > Virtual Servers > your virtual server > Server Properties**.
2. In the left menu, select **Monitoring Policy**.
3. Click **Lock**.
4. In the **Layer 3 Monitoring** section, specify the IP address monitoring policy. For more information, see [Layer 3 Monitoring](#).
5. In the **Monitored IPs I / II** tables, add the IP addresses that must be reachable via the ICMP protocol by the system that is hosting the server.
6. In the **Layer 2 Monitoring** section, specify the interface monitoring policy. For more information, see [Layer 2 Monitoring](#).
7. In the **Monitored Interfaces I / II** tables, add the physical interfaces that must have a link in order for the server to stay up.
8. Click **Send Changes** and **Activate**.

Configure Custom Scripts

Configure custom scripts for use with your monitoring policies. These scripts are run after the server starts or before the server shuts down due to unreachable IP addresses or interfaces.

Do not use `phionctrl` in your custom scripts; this might cause a deadlock.

1. Go to **CONFIGURATION > Configuration Tree > Box > Virtual Servers > your virtual server > Server Properties**.
2. In the left menu, select **Custom Scripts**.
3. Click **Lock**.
4. In the **Start** and **Stop Script** fields, enter the commands that should be executed when the server is started up or shut down (7-bit ASCII characters and standard Bash version 2-compliant).
5. Click **Send Changes** and **Activate**.

Figures

1. pol1.png
2. pol2.png

© 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.