

High Availability in Azure

<https://campus.barracuda.com/doc/79463116/>

Deploying a Barracuda CloudGen Firewall HA cluster in the Microsoft Azure Public Cloud requires a custom deployment and configuration to integrate the CloudGen Firewall with the Azure networking stack. Since the IP addresses are fixed and non-transferable between VMs, incoming traffic is handled and delivered via the Azure Load Balancer. This load balancer polls a service that is reachable through the Forwarding Firewall service on the active firewall. When the service fails over to the other unit, the load balancer will follow with a latency of a couple of seconds. For backend servers to be able to send traffic through the currently active unit, there are two high availability methods available for Azure:

- **Cloud Integration** - With this method, the Firewall can directly manipulate the Azure routing table so that routing entries always point to the active unit of the HA cluster. Due to the limitation of Azure networking, all active sessions will time out whenever a failover occurs.

or

- **Standard Load Balancer** - With this method, a new type of load balancer is used in Azure to be the destination for the route tables. The load balancer will monitor for the active firewall and direct traffic on all ports to the active unit.

Barracuda Firewall Control Centers in Azure do not support high availability configurations.

Azure Route Tables and Standard Load Balancers

The fastest failover method available within Azure is to use the new Standard Load Balancer product from Microsoft. When using a Standard Load Balancer, the route tables will point at an internal load balancer's IP. The load balancer actively probes the firewalls to determine which one is active. Using the new rule type **HA Ports**, it will then redirect all traffic to the active unit. This method is stateful and fails over as quickly as the load balancer probes can indicate failure.

For more information, see: [How to Configure a High Availability Cluster in Azure with the Standard Load Balancer](#).

You cannot directly migrate between a Cloud Integration design and a Standard Load Balancer design because the load balancers and any public IPs used need to be built using a new Standard type. It is not possible to combine the old IPs, which are considered Basic, with any Standard components.

Standard load balancer designs require a Network Security Group (NSG) to be present in order to pass any traffic. The NSG can permit Any inbound and Any outbound to allow the firewall to filter instead.

Choosing a Failover Method

Cloud Integration	Standard Load Balancers
No additional costs	Charges per rule and traffic volume
Slow failover as API calls	Fast failover as LB monitors
Non-stateful due to routes being rewritten	Stateful failover
Poor for multi-VNET designs	Best for multi-VNET designs

Deploy a High Availability Cluster

Create a high availability cluster by deploying two CloudGen Firewall VMs in the same subnet and availability zone. Incoming traffic is then forwarded to the active firewall by an Azure Load Balancer. User-defined routing and the rewriting of the routes by the active firewall ensure that the backend VMs always use the active firewall as the gateway device.

For more information, see [How to Configure a High Availability Cluster in Azure using PowerShell and ARM](#).

Azure Route Table Rewriting with Azure Cloud Integration

User-defined routing is limited to one VM as the gateway device when creating a route. When you are using a high availability cluster as the gateway, the VM that is used changes when the service fails over. Both firewalls are configured to access the Azure fabric and reconfigure the routing table when a failover occurs. For the firewall to rewrite the routes, you must configure Azure cloud integration.

For more information, see [How to Configure Azure Cloud Integration using ARM](#).

Example Video

Watch the following video to see a high availability firewall cluster using SD-WAN for a hybrid cloud setup.

Barracuda Azure Solutions Training

NextGen Firewall Hybrid Cloud Setup

Felix Bueltmann (TME NextGen/Cloud)



© Barracuda Networks Inc., 2019 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.