



# High Availability in Azure

Deploying a Barracuda NextGen Firewall F-Series HA cluster in the Microsoft Azure Public Cloud requires a custom deployment and configuration to integrate the F-Series 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 either a service running on the virtual server of the F-Series Firewall or a service that is reachable only through the forwarding firewall service on the active F-Series Firewall. When the virtual server 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, the F-Series 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. Barracuda NextGen Control Centers in Azure do not support high availability configurations.

## Azure Resource Manager

### Deploy a High Availability Cluster

Create a high availability cluster by deploying two F-Series 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 ensures 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

Azure 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 virtual server 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 an high availability firewall cluster using Traffic Intelligence for a hybrid cloud setup.

## Azure Service Manager

### Deploy a High Availability Cluster via Web UI

For legacy Azure deployments using Azure Service Manager, deploying the F-Series Firewall cluster via the Azure Web UI is easier for first-time users. Deploying the HA cluster via the web UI limits the use of advanced Azure networking features only available via PowerShell.

For more information, see [How to Configure a High Availability Cluster in Azure via Web Portal and ASM](#).

### Deploy a High Availability Cluster via PowerShell

Deploying via PowerShell scripts grants you more flexibility and the use of features that are only available via PowerShell.

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



### **Azure Route Table Rewriting with Azure Cloud Integration**

Azure 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 virtual server 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 ASM](#).

