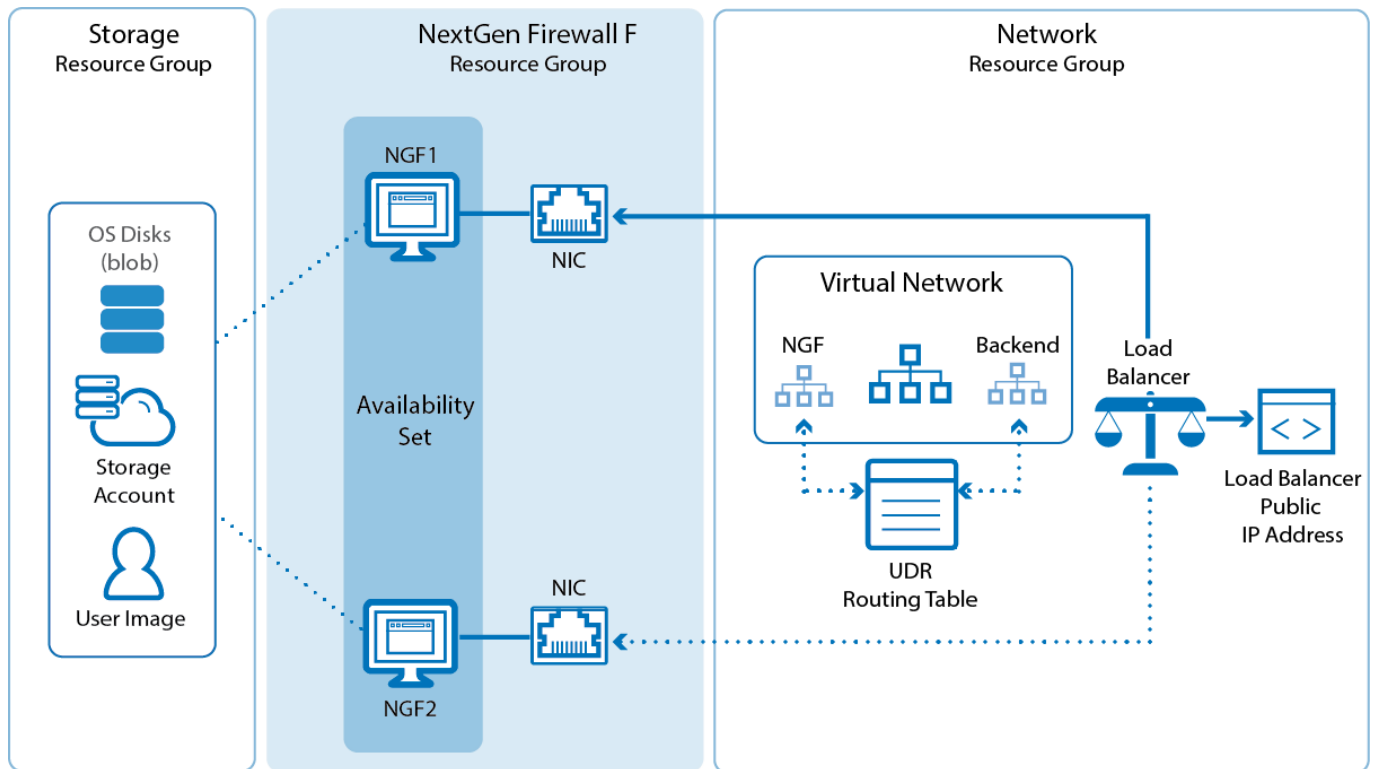


How to Configure a High Availability Cluster in Azure using PowerShell and ARM

<https://campus.barracuda.com/doc/46897612/>

Configure a high availability cluster to ensure that the services running on the Barracuda NextGen Firewall F-Series VMs are always available even if one unit is unavailable due to maintenance or a hardware issue. To be able to configure an HA cluster, both firewalls VMs must be deployed to the same subnet and be placed in an Availability Set. This ensures that the VMs are placed in different fault and update domains inside the Azure datacenter. Incoming connections are forwarded to the active firewall by the Azure Load balancer. The load balancer actively monitors the services on the firewall and, when an HA failover takes place, redirects the traffic to the other, now-active firewall. You must create load balancer rules and health probes for each service for the load balancer to know which ports to forward and how to monitor them. The load balancer does not failover immediately after the virtual server has failed over, since it requires at least two probes to fail before reacting. Combined with the minimum poll time of 5 seconds, this means that failover will take at least 10 seconds during which no traffic can be forwarded.

The existing session will time out after a failover since Azure does not support floating IP addresses; the connection to the backend server VMs will thus use a different source IP address. The backend VMs are configured to use the firewall as the default gateway and, if needed, access control between the backend subnets using Azure User Defined Routing. Because only one IP address can be configured as the destination, the F-Series Firewall connects to the Azure fabric and rewrites the route table to use the now-active firewall instead of the failed firewall VM. Now, the backend VMs can connect via the active firewall to the Internet.



In this article:

Before You Begin

- Install Azure PowerShell 1.1.0 or higher.

Step 1. Deploy two NextGen F-Series Firewall VMs

To configure an HA cluster, deploy two F-Series VMs. The Public IPs attached to the NICs are removed after configuring client-to-site VPN access via the load balancer. To be able to use them in an HA cluster, the deployment must meet the following requirements:

- Single Network Interfaces must be used.
- Both VMs must be in the same subnet of the virtual network.
- Both VMs must be in the same subnet.
- Static private (internal) IP addresses must be used.
- The same instance size for both VMs must be used.
- Both firewalls must be the same Barracuda NextGen Firewall F for Azure model.

- Both VMs must be deployed in one Availability Set.

For more information, see [How to Deploy an F-Series Firewall in Microsoft Azure using Azure Portal and ARM](#) or [How to Deploy an F-Series Firewall in Microsoft Azure using PowerShell and ARM](#).





Step 2. Change the Firewall Network Configuration to Use the Static Private IP Addresses

On both firewall VMs, change the network configuration to use a static network interface. Use the static private IP address you assigned to the NIC during deployment.

Step 2.1 Reconfigure the Network Interface

Change the network interface type from dynamic to static.


1. Go to **CONFIGURATION > Configuration Tree > Box > Network**.
2. In the left menu, click on **xDSL/DHCP/ISDN**.
3. Click **Lock**.
4. Delete the **DHCP01** entry in the **DHCP Links** list.
5. Select **No** from the **DHCP Enabled** drop-down list.
6. Click **Send Changes**.
7. In the left menu, click on **IP Configuration**.
8. In the **Management IP and Network** section in the **Interface Name** line, clear the **Other** check box.
9. Select **eth0** from the **Interface Name** list.
10. Enter the static internal IP address from [Step 1](#) as the **Management IP (MIP)**.
E.g., 10.0.20.6

Management IP and Network	
Interface Name	<input checked="" type="checkbox"/> eth0 <input type="checkbox"/> Other 
Management IP (MIP)	<input checked="" type="checkbox"/> 10.0.20.6   
Associated Netmask	24-Bit 
Responds to Ping	yes 
Use for NTPd	yes 
Advertise Route	no 

Step 2.2 Create the Default Route

Add the default route. The default gateway in Azure subnets is always the first IP in the subnet. E.g., 10.0.20.1 if the subnet is 10.0.20.0/24

1. In the left menu, click on **Routing**.
2. Click **+** in the **Routes** table and configure the following settings:
 - **Target Network Address** - Enter 0.0.0.0/0
 - **Route Type** - Select **gateway**.
 - **Gateway** - Enter the first IP address of the subnet the firewalls reside in. E.g., 10.0.20.1 if the IP addresses of the firewalls are 10.0.20.6 and 10.0.20.7
 - **Trust Level** - Select **Unclassified**.



Route Configuration	
Target Network Address	0.0.0.0/0
Route Type	gateway
Interface Name	<input type="text"/> Other
Gateway	10.0.20.1
Route Metric	<input type="text"/>
Source Address	<input type="text"/>
Trust Level	Unclassified

3. Click **OK**.
4. Click **Send Changes** and **Activate**.

Step 2.3 Activate the Network Changes

Activate the changes to the network configuration.

1. Go to **CONTROL > Box**.
2. In the **Network** section of the left menu, click on **Activate new network configuration**.
3. Click **Failsafe**.

Open the **CONTROL > Network** page. Your interface and IP address are now static.

Step 3. Configure an HA Cluster on the F-Series Firewall VMs

Configure the two firewalls to synchronize session and configuration information. Because Azure does not support floating IP addresses, you must configure all services on the virtual server to listen on a loopback address (127.0.0.X). Use **Application Redirect** access rules to redirect incoming traffic from the eth0 interface to the services. Use the the internal IP address of the primary and secondary firewall as the destination of the rule to ensure that it matches without regard to which firewall VM the virtual server is currently running on.

For more information, see [High Availability](#) and [How to Set Up a High Availability Cluster](#).

Step 4. Configure the Azure Load Balancer

Deploy an Azure Load Balancer to monitor and forward incoming traffic to the active firewall. You must add a load balancer rule for each service with at least one health probe. The probe determines how the health of the service is checked. Load balancer rules are required for the following services:

- **TINA VPN on TCP or UDP 691** – If both TCP and UDP are required you must map one protocol to a different internal port such as 693 and then use an Application redirect rule to forward the traffic to the VPN service on port 691. Probe for TCP on port 691.
- **SSH on TCP 22** – This load balancer rule will allow to connect to the active firewall with ssh. The secondary firewall can be accessed via the active firewall.

For more information, see [How to Configure Azure Load Balancer for HA Clusters using PowerShell and ARM](#).

Step 5. Configure User Defined Routing in Azure

Configure UDR for the backend VMs to use the firewall as their default gateways for all connections to the Internet and/or between backend subnets.

For more information, see [How to Configure Azure Route Tables \(UDR\) using PowerShell and ARM](#).

Step 6. Configure HA UDR on the F-Series Firewall VMs

Configure both the primary and secondary firewalls to update the UDR routing table. In case of a failover, an update to the Azure Route table is triggered so that all backend VMs using the HA cluster as a gateway use the active firewall.

For more information, see [How to Configure Azure Route Table Rewriting for HA Clusters using ARM](#).

Step 7. Configure a Client-to-Site VPN for Management Access

Configure a TINA client-to-site VPN that will be used for management access. Connect via the load

balancer public IP address.

For more information, see [How to Configure a Client-to-Site TINA VPN with Personal Licenses](#) or [How to Configure a Client-to-Site TINA VPN with Client Certificate Authentication](#).

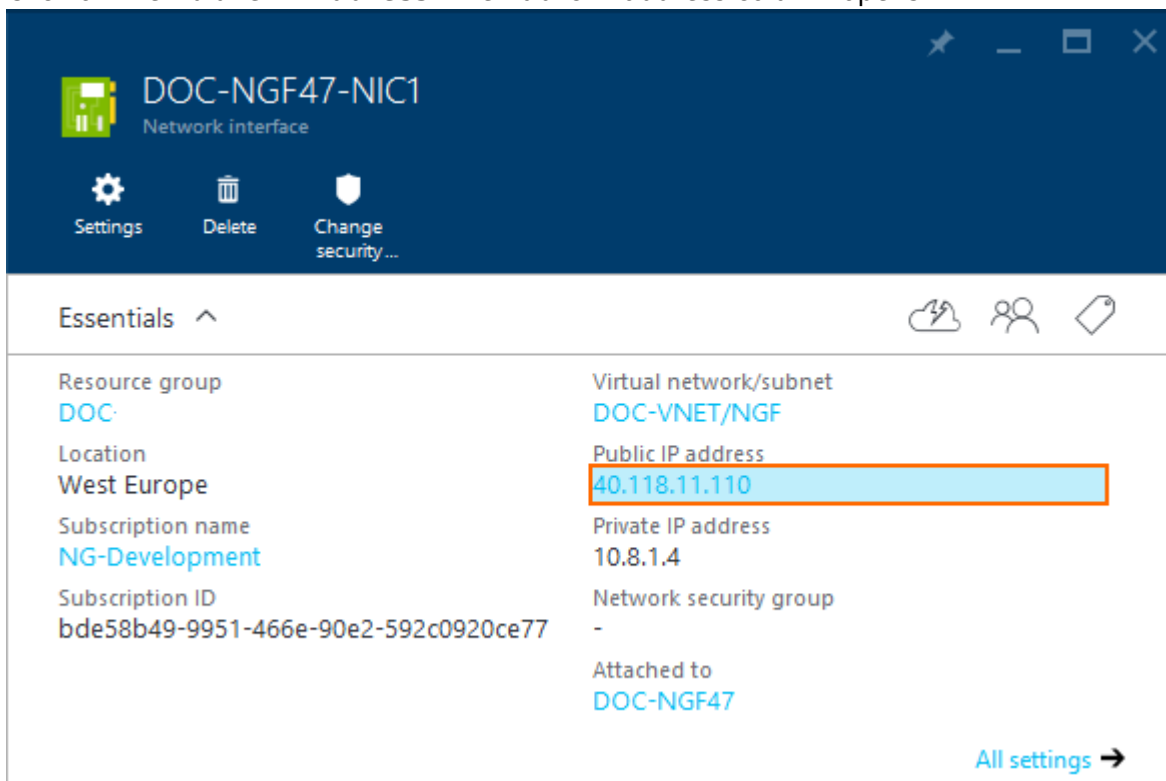
Step 8. Disassociate the Public IP Addresses

When both a loadbalancer and a public IP are available for the firewall VM the public IP is used as the default source IP address for the VM. This would mean that outgoing connections would use different source IP addresses depending on which firewall is active.

Using the Azure Web Portal

For each firewall VM remove the Public IP address from the network interface.

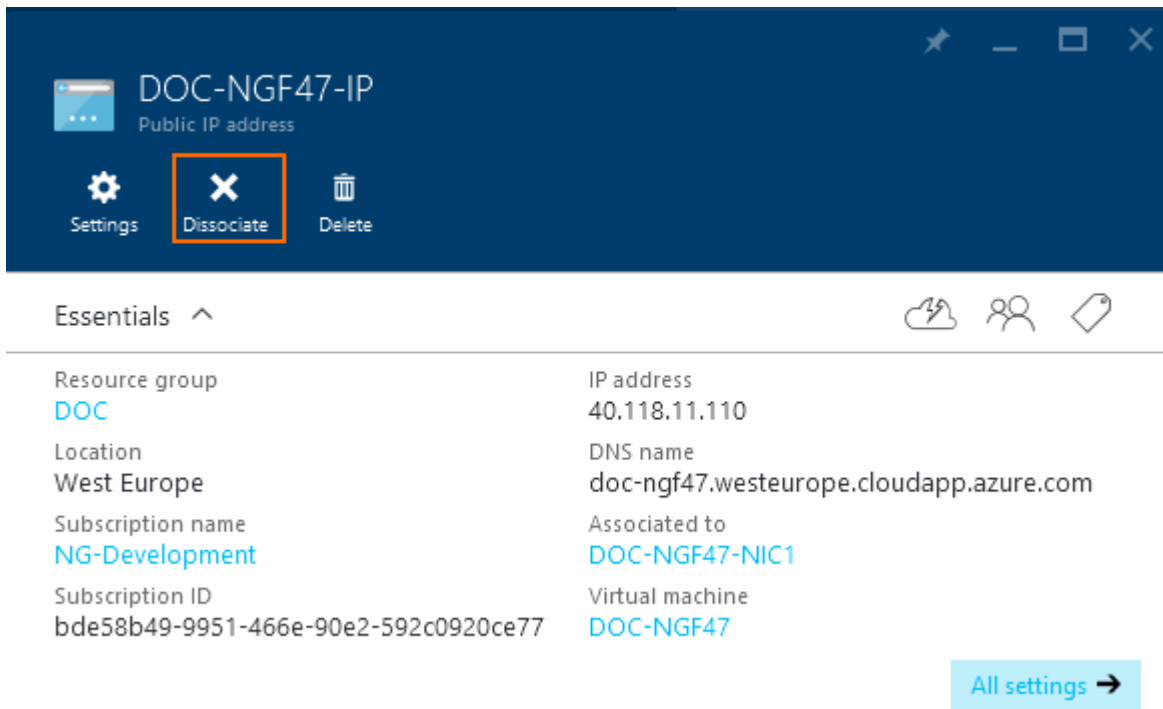
1. Go to <https://portal.azure.com>.
2. Locate the Network Interface attached to your primary firewall VM.
3. Click on the **Public IP Address**. The Public IP address column opens.



The screenshot shows the Azure portal interface for a network interface named 'DOC-NGF47-NIC1'. The interface is located in the 'Essentials' section. The public IP address '40.118.11.110' is highlighted with a red box. The private IP address is '10.8.1.4'. The network interface is attached to the virtual network/subnet 'DOC-VNET/NGF'.


Property	Value
Resource group	DOC
Location	West Europe
Subscription name	NG-Development
Subscription ID	bde58b49-9951-466e-90e2-592c0920ce77
Virtual network/subnet	DOC-VNET/NGF
Public IP address	40.118.11.110
Private IP address	10.8.1.4
Network security group	-
Attached to	DOC-NGF47

4. Click **Disassociate**.



5. Repeat for the secondary firewall VM.

You can now connect to the client-to-site VPN and manage both Barracuda NextGen Firewall F VMs via the internal IP addresses. Go to **CONTROL > Network > Azure UDR** and verify that all routes using the firewall VM are green.



Figures

1. azure_arm_ha_diagram.png
2. AzureHA08.png
3. Azure_default_route.png
4. AzureHA00.png
5. AzureHA01.png
6. AzureHA02.png

© 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.