

How to Configure a High Availability Cluster in Azure via PowerShell and ASM

<https://campus.barracuda.com/doc/48203164/>

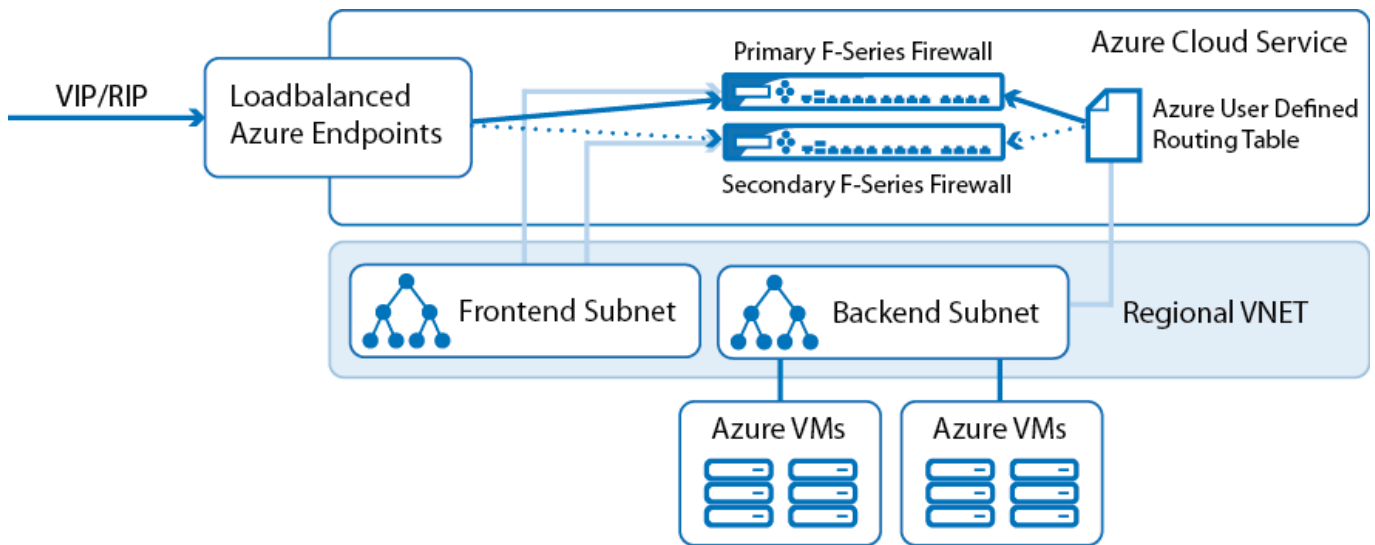
To safeguard against hardware and software failures in the Azure cloud, use a high availability (HA) setup. Most advanced networking features, such as static internal IP addresses or a reserved IP address for the Cloud Service, can only be deployed via Azure PowerShell. You can either enter the commands directly into the Azure PowerShell or combine the commandlets to a custom deployment script.

Both Barracuda NextGen Firewall F-Series VMs must be deployed in the same Azure Availability Set, so that the VMs run in different fault domains in the Azure datacenter. Both NextGen Firewalls are connected to the same Azure virtual network and use static internal IP addresses (DIPs). An Azure load-balanced endpoint (level 4 load balancer) can be used to offer TCP- and UDP-based services on the VIP.

Connecting to Services and Managing the HA cluster in the Azure cloud:

1. **Accessing Services in Azure/ on the NextGen Firewall F-Series** – Create a Load-balanced Endpoint for each service accessed on or behind the Barracuda NextGen Firewall F-Series in Azure.
2. **Management Access** – If you are not using a Barracuda NextGen Control Center to manage your NextGen Firewall F-Series use the following solution to be able to access both VMs with NG Admin:
 - Create an Endpoint on port TCP/807 to manage the primary NextGen Firewall F-Series.
 - Configure a Client-to-Site VPN. You can now reach the static internal IP address of the secondary NextGen Firewall F-Series through the Client-to-Site VPN.

Azure (load-balanced) Endpoints can be used only for TCP/UDP-based services. All other IP protocols (ICMP, ESP,...) are blocked.



Before you begin

Install the latest Windows PowerShell for Azure. (<http://azure.microsoft.com/en-us/downloads/>)

Step 1. Create an Azure Virtual Network

You must use a Regional VNET to deploy the Barracuda NextGen Firewall F-Series. Older VNETs using Affinity Groups do not support the necessary Azure networking features. Configuration information of the VNet is stored in an XML file and then deployed in the Azure cloud via PowerShell commandlet. The following is an example `vmnet.xml` with 2 subnets:

```
<NetworkConfiguration
  xmlns="http://schemas.microsoft.com/ServiceHosting/2011/07/NetworkC
onfiguration">
  <VirtualNetworkConfiguration>
    <VirtualNetworkSites>
      <VirtualNetworkSite name="NEVNET"
Location="North Europe">
        <AddressSpace>
          <AddressPrefix>10.0.0.0/16</AddressPrefix>
        </AddressSpace>
        <Subnets>
          <Subnet name="Frontend">
            <AddressPrefix>10.0.30.0/24</AddressPrefix>
          </Subnet>

```

```

    <Subnets>;
    <Subnet name="Backend">;
      <AddressPrefix>10.0.31.0/24</AddressPrefix>;
    </Subnet>;
  </Subnets>;
  &nbsp;</VirtualNetworkSite>;
</VirtualNetworkSites>;
</VirtualNetworkConfiguration>;
</NetworkConfiguration>;

```

1. Open the Azure PowerShell.
2. If VNets already exist, export the existing Virtual Networks to an XML file.
`Get-AzureVNetConfig -ExportToFile c:\azure\vmnet.xml`
3. Edit the vnet.xml file and enter the configuration for your **VIRTUALNETWORKSITE**. Use the example file above as a guideline. If you are using multiple network interfaces, create one subnet per network interface.
4. Upload the VNet configuration file:
`Set-AzureVNetConfig -ConfigurationPath c:\azure\modified_vmnet.xml`

The virtual network is now listed in **VIRTUAL NETWORKS** in the web UI, via PowerShell:

`Get-AzureVNetSite -VNetName "YOUR VNET NAME"`

```

PS C:\> Get-AzureVNetSite -VNetName NEUNET
VERBOSE: 09:12:25 - Begin Operation: Get-AzureVNetSite
VERBOSE: 09:12:29 - Completed Operation: Get-AzureVNetSite

AddressSpacePrefixes : <10.0.0.0/16>
AffinityGroup         :
DnsServers            : <>
GatewayProfile        :
GatewaySites          :
Id                    : ee12e308-8ef1-4906-b0e9-8d40f4093cf0
InUse                  : False
Label                 :
Name                  : NEUNET
State                  : Created
Subnets              : <Backend, Frontend, Subnet3, Subnet4>
OperationDescription  : Get-AzureVNetSite
OperationId           : 5021fd31-1cb4-3c18-bf86-ba0b2cec25d9
OperationStatus       : Succeeded
PS C:\>

```

Step 2. (optional) Use a reserved IP for the Azure cloud service

To avoid the difficulty of a changing public IP address when redeploying your Cloud Service, you can reserve a public IP address and assign it to your cloud service on deployment. This IP address persists even when the cloud service is deleted.

Create a Reserved IP address (RIP).

```
New-AzureReservedIP -ReservedIPName "RIP NAME" -Label "NextGen Firewall F-  
Series IP" -Location "YOUR LOCATION"
```

Enter the name of the Reserved IP address in the example deployment script below to always use the same public IP address for your Cloud Service.

Step 3. Create a storage account

1. Create a Storage Account and set it as the default storage account.


```
New-AzureStorageAccount -StorageAccountName "STORAGEACCOUNT NAME" -  
Location "YOUR LOCATION"
```

2. Use the storage account as the default storage account for this Azure subscription.

```
Set-AzureSubscription -SubscriptionName "YOUR AZURE SUBSCRIPTION NAME" -  
CurrentStorageAccountName "STORAGEACCOUNT NAME"
```

3. Verify that you are using the correct storage account:

```
Get-AzureSubscription
```



```
PS C:\> Get-AzureSubscription  
SubscriptionId      : ee21fa5b-30b7-45d1-81c9-919085006474  
SubscriptionName    : Pay-As-You-Go  
Environment         : AzureCloud  
SupportedModes      : AzureServiceManagement  
DefaultAccount      : EFB956162E48C23A67E69356463356C42FB40729  
Accounts            : <EFB956162E48C23A67E69356463356C42FB40729>  
IsDefault           : True  
IsCurrent            : True  
CurrentStorageAccountName : docstorage02  
PS C:\> _
```

Step 4. Deploy two Barracuda NextGen Firewall F-Series VMs

Deploy the primary and secondary Barracuda NextGen Firewalls using PowerShell. You can change the example deployment script below to fit your needs and deploy both Barracuda NextGen Firewalls at the same time. Deploy the two NextGen Firewalls as to use:

- Static internal IP addresses in the Frontend subnet
- Assign both to the same Availability Set.
- Install and activate the license on the primary NextGen Firewall F-Series. The secondary NextGen Firewall F-Series receives the license through the HA-sync when the HA cluster is paired. For more information, see [How to Activate and License a Stand-alone Hardware](#)

[NextGen F-Series or Control Center Appliance.](#)

This script will deploy two Barracuda NextGen Firewalls with an Endpoint for NG Admin for the primary NextGen Firewall F-Series, and a Load Balanced Endpoint for the secondary NextGen Firewall F-Series.

```
#If needed import Azure PSD file #Import-Module "C:\Program Files
(x86)\Microsoft SDKs\Azure\PowerShell\ServiceManagement\Azure\Azure.psd" #
Use Default System Proxy [System.Net.WebRequest]::DefaultWebProxy.Credentials
= [System.Net.CredentialCache]::DefaultCredentials $HA1vmname = "NGFW1"
$HA2vmname = "NGFW2" $RootPassword = "B@rr@cu@2015" $instanceSize = "Large"
$cloudService = "NGFWCS" $Location = "North Europe" $storageAccount
="YOURSTORAGEACCOUNT" $reservedIPname = "" $VNetName = "NEVNET" $Subnet1 =
"Frontend" $Subnet2 = "Backend" $HA1NIC1IP = "10.0.30.20" $HA1NIC2IP =
"10.0.31.31" $HA2NIC1IP = "10.0.30.21" $HA2NIC2IP = "10.0.31.32" #Enter a VM
Image name below to use a custom image. If left empty the latest image from
the Azure Marketplace is used. $image = "" $availabilitySetName
="BarracudaNGAVSet" $azureSubscriptionName = "Pay-As-You-Go" function
AskYesNo( $title, $question, $YesInfo, $NoInfo ) { $yes = New-Object
System.Management.Automation.Host.ChoiceDescription "&Yes", $YesInfo $no =
New-Object System.Management.Automation.Host.ChoiceDescription "&No", $NoInfo
$options = [System.Management.Automation.Host.ChoiceDescription[]]($yes, $no)
$result = $host.ui.PromptForChoice($title, "Deploy the Barracuada NextGen
Firewall F-Series(s) as shown?", $options, 0) return $result } # Show summary
of Barracuda NextGen Firewall F-Series that is to be deployed and then ask
for permission Write-Host -NoNewLine "This script will create a " Write-Host
-NoNewLine -ForegroundColor yellow "dual-NIC Barracuda NextGen Firewall F-
Series HA Cluster" Write-Host " in Azure" Write-Host "" Write-Host -NoNewLine
"VNET name: " Write-Host -ForegroundColor yellow $VNetName Write-Host -
NoNewLine "Azure DC: " Write-Host -ForegroundColor yellow $Location Write-
Host "Primary NextGen Firewall F-Series NIC 1: " Write-Host -NoNewLine -
ForegroundColor yellow "$HA1NIC1IP in $Subnet1" Write-Host " (management)"
Write-Host -NoNewLine "NIC 2: " Write-Host -ForegroundColor yellow
"$HA1NIC2IP in $Subnet2" Write-Host "Secondary NextGen Firewall F-Series NIC
1: " Write-Host -NoNewLine -ForegroundColor yellow "$HA2NIC1IP in $Subnet1"
Write-Host " (management)" Write-Host -NoNewLine "NIC 2: " Write-Host -
ForegroundColor yellow "$HA2NIC2IP in $Subnet2" $yesorno = AskYesNo 'Do you
want to continue?' $warn 'will abort script' 'will use existing VNet' switch
( $yesorno ) { 0 { "OK! Creating two new Barracuda NextGen Firewall F-Series
VMs." } 1 { "Got it :( Please correct variable values in script and rerun."
return } } # Create storage if it doesn't exist yet if(!(Test-AzureName -
Storage $storageAccount)) { Write-Host "Creating Storage Account
$storageAccount in $Location" New-AzureStorageAccount -StorageAccountName
$storageAccount -Location $Location } # Reserved IP for the Cloud Service
#$reservedIP = New-AzureReservedIP -ReservedIPName $reservedIPname -Label
$reservedIPname -Location $location if ($reservedIPName -ne "") { $reservedIP
```

```
= Get-AzureReservedIP -ReservedIPName $reservedIPName Write-Host "Using Existing Reserved IP!" } # Set storage account as default storage Set-AzureSubscription -SubscriptionName $azureSubscriptionName -CurrentStorageAccountName $storageAccount # If no explicit image is defined get the latest Barracuda NextGen Firewall F-Series Azure Image available in the Azure Marketplace if ( $image -eq "" ) { $image = Get-AzureVMImage | where { $_.ImageFamily -Match "Barracuda NextGen Firewall F-Series*" } | where { $_.Location.Split(";") -contains $Location } | sort PublishedDate -Descending | select -ExpandProperty ImageName -First 1 Write-Host "Using Image from Azure Marketplace..." } # Create Azure VMs $halvm = New-AzureVMConfig -Name $halvmname -InstanceSize $instanceSize -Image $image -AvailabilitySetName $availabilitySetName Add-AzureProvisioningConfig -Linux -LinuxUser "azureuser" -Password $RootPassword -VM $halvm $ha2vm = New-AzureVMConfig -Name $ha2vmname -InstanceSize $instanceSize -Image $image -AvailabilitySetName $availabilitySetName Add-AzureProvisioningConfig -Linux -LinuxUser "azureuser" -Password $RootPassword -VM $ha2vm -NoSSHEndpoint # Add Endpoints for 1st NIC of the primary Barracuda NextGen Firewall F-Series Write-Host "Adding Endpoints for primary Barracuda NextGen Firewall F-Series ..." Add-AzureEndpoint -Protocol tcp -LocalPort 22 -PublicPort 222 -Name "SSH" -VM $halvm Add-AzureEndpoint -Protocol tcp -LocalPort 807 -PublicPort 807 -Name "MGMT" -VM $halvm # Add Endpoints for 1st NIC of the secondary Barracuda NextGen Firewall F-Series Write-Host "Added Endpoints for secondary Barracuda NextGen Firewall F-Series ..." Add-AzureEndpoint -Protocol tcp -LocalPort 22 -PublicPort 22 -Name "SSH" -VM $ha2vm #Add-AzureEndpoint -Protocol tcp -LocalPort 807 -PublicPort 807 -Name "MGMT" -VM $ha2vm1 # Create Loadbalanced Endpoint for TCP TINA VPN Add-AzureEndpoint -Name "HA1VPNTCP" -Protocol tcp -LocalPort 691 -PublicPort 691 -LBSetName "LBVPNTCP" -DefaultProbe -VM $halvm Add-AzureEndpoint -Name "HA2VPNTCP" -Protocol tcp -LocalPort 691 -PublicPort 691 -LBSetName "LBVPNTCP" -DefaultProbe -VM $ha2vm # Create Loadbalanced Endpoint for UDP TINA VPN Add-AzureEndpoint -Protocol udp -LocalPort 691 -PublicPort 691 -Name "HA1VPNUDP" -LBSetName "LBVPNUDP" -DefaultProbe -VM $halvm Add-AzureEndpoint -Protocol udp -LocalPort 691 -PublicPort 691 -Name "HA2VPNUDP" -LBSetName "LBVPNUDP" -DefaultProbe -VM $ha2vm # Create Loadbalanced Endpoint for HTTP Traffic. Set the probes to monitor a service on the NextGen Firewall F-Series #Add-AzureEndpoint -Protocol tcp -LocalPort 80 -PublicPort 80 -Name "HA1VPNHTTP" -LBSetName "LBHTTP" -ProbePort 691 -ProbeProtocol tcp -VM $halvm #Add-AzureEndpoint -Protocol tcp -LocalPort 80 -PublicPort 80 -Name "HA2VPNHTTP" -LBSetName "LBHTTP" -ProbePort 691 -ProbeProtocol tcp -VM $ha2vm # Create Loadbalanced Endpoint for HTTP Traffic. Set the probes to monitor a service on the NextGen Firewall F-Series #Add-AzureEndpoint -Protocol tcp -LocalPort 443 -PublicPort 443 -Name "HA1VPNHTTPS" -LBSetName "LBHTTPS" -ProbePort 691 -ProbeProtocol tcp -VM $halvm #Add-AzureEndpoint -Protocol tcp -LocalPort 443 -PublicPort 443 -Name "HA2VPNHTTPS" -LBSetName "LBHTTPS" -ProbePort 691 -ProbeProtocol tcp -VM $ha2vm # Define Subnet and static IP Address for 1st NIC Set-AzureSubnet -SubnetName $Subnet1 -VM $halvm Set-AzureStaticVNetIP -IPAddress
```

```
$HA1NIC1IP -VM $ha1vm Set-AzureSubnet -SubnetName $Subnet1 -VM $ha2vm Set-
AzureStaticVNetIP -IPAddress $HA2NIC1IP -VM $ha2vm Write-Host "Configured
First NIC..." # Add Additional NICs Add-AzureNetworkInterfaceConfig -Name
"HA1NIC2" -SubnetName $Subnet2 -StaticVNetIPAddress $HA1NIC2IP -VM $ha1vm
Add-AzureNetworkInterfaceConfig -Name "HA2NIC2" -SubnetName $Subnet2 -
StaticVNetIPAddress $HA2NIC2IP -VM $ha2vm Write-Host "Added second NIC..."
Write-Host "Starting VM deployments. This may take a while..." if
($reservedIPName -eq "") { # Create NG without reservedIP Write-Host
"Creating primary Barracuda NextGen Firewall F-Series." New-AzureVM -
ServiceName $cloudService -VM $ha1vm -Location $Location -VNetName $VNetName
Write-Host "Creating secondary Barracuda NextGen Firewall F-Series." New-
AzureVM -ServiceName $cloudService -VM $ha2vm } else { # Create NG With
Reserved IP address Write-Host "Creating primary Barracuda NextGen Firewall
F-Series using Reserved IP address $reservedIPName" New-AzureVM -ServiceName
$cloudService -VM $ha1vm -ReservedIPName $reservedIPName -Location $Location
-VNetName $VNetName Write-Host "Creating secondary Barracuda NextGen Firewall
F-Series." New-AzureVM -ServiceName $cloudService -VM $ha2vm } Write-Host
"Script is done. Creating the Virtual Machines can take a while. Have a cup
of coffee! Use Barracuda NextGen Admin to login to
$cloudService.cloudapp.net: user: root, password: $RootPassword)"
```

Step 5. Create additional load balanced endpoints

The example script above creates a Load Balanced Endpoint for a TINA VPN using tcp by default. When creating Endpoints used for services that are not hosted on the Barracuda NextGen Firewall F-Series, you must configure the probe port and protocol to probe a service running on the NextGen Firewall F-Series to allow the load balancer to detect an HA failover. For example, when configuring a load balanced Endpoint for a web server behind the NextGen Firewall F-Series, configure the probe port and protocol to monitor the VPN service running on port 691 instead.

To add an additional Load Balanced Endpoint after the Barracuda NextGen Firewalls have been deployed, enter:

```
Get-AzureVM -ServiceName YOUR_CLOUD_SERVICE -Name
YOUR_PRIMARY_BARRACUDA_NG_FIREWALL | Add-AzureEndpoint -Protocol tcp -
LocalPort 443 -PublicPort 443 -Name "HA1HTTPS" -LBSetName "LBNAME" -
DefaultProbe | Update-AzureVM Get-AzureVM -ServiceName YOUR_CLOUD_SERVICE -
Name YOUR_SECONDARY_BARRACUDA_NG_FIREWALL | Add-AzureEndpoint -Protocol tcp -
LocalPort 443 -PublicPort 443 -Name "HA2HTTPS" -LBSetName "LBNAME" -
DefaultProbe | Update-AzureVM
```


Step 6. Change the firewall network configuration to use the static internal IP addresses

Change the network configuration of the primary and secondary Barracuda NextGen Firewall F-Series to use a static network interface that you previously configured during deployment of the virtual machines.

Step 6.1 Reconfigure the network interface

Change the network interface type from dynamic to static.

1. Log into the primary Barracuda NextGen Firewall F-Series.
2. Go to **CONFIGURATION > Configuration Tree > Box > Network**.
3. In the left menu, click on **xDSL/DHCP/ISDN**.
4. Click **Lock**.
5. Delete the **DHCP01** entry in the **DHCP Links** list.
6. Select **No** from the **DHCP Enabled** drop-down list



DHCP Client Setup


DHCP Enabled: yes

DHCP Links

Name	Link Active	Standby Mode
DHCP01	no	

Use these links to connect to a cable modem or a DSL line via an external DSL router.

7. Click **Send Changes**.
8. In the left menu, click on **IP Configuration**.
9. In the **Management IP and Network** section in the **Interface Name** line, disable the **Other** checkbox.
10. Select **eth0** from the **Interface Name** list.
11. Enter the static internal IP address used in the deployment script in Step 4 as the **Management IP (MIP)**. E.g., 10.0.30.20



Management IP and Network

Interface Name: eth0 Other

Management IP (MIP): 10.0.30.20

Associated Netmask: 24-Bit


Responds to Ping: yes

Use for NTPd: yes

Step 6.3 Create the default route

Add the default route.

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 Barracuda NextGen Firewalls reside in. E.g., `10.0.30.1` if the IP addresses of the Barracuda NextGen Firewalls are `10.0.30.20` and `10.0.30.21`
 - **Trust Level** – Select **Unclassified**.



Route Configuration

Target Network Address	<input type="text" value="0.0.0.0/0"/>
Route Type	<input type="text" value="gateway"/>
Interface Name	<input type="text" value=""/> <input type="checkbox"/> Other
Gateway	<input type="text" value="10.0.30.1"/>
Route Metric	<input type="text" value=""/>
Source Address	<input type="text" value=""/>
Trust Level	<input type="text" value="Unclassified"/>

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

Step 6.4 Add and configure the second network interface

You must configure the second NIC on each NextGen Firewall F-Series. The second interface can be configured to be static or dynamic. If you are using a static network interface you must use the IP address you assigned to the second NIC of the VM during deployment. Add the second IP address to the Additional Local IPs as it can only be used as a box layer IP address. This address will not be switched over to the secondary NextGen Firewall F-Series during a failover.

1. In the left menu, click on **Interfaces**.
2. Click **Lock**.
3. Double click on the **netsvc** entry in the **Network Interface Cards** list. The **Network Interface Cards** window opens.
4. Set **Number of Interfaces** to **2**.
5. Click **OK**.
6. Click **Send Changes**.
7. In the left menu, click on **IP Configuration**.
8. In the **Additional Local IPs** section, click **+** to add a **IP Address Configuration**. The **IP Address Configuration** window opens.
9. Enter a **Name**.

10. Click **OK**.
11. Enter the following settings for the **IP Address Configuration**:
 - **Interface Name** - Select **eth1**.
 - **IP Address** - Enter the IP address you assigned to this interface during deployment.
 - **Associated Netmask** - Select the netmask for this subnet.
 - **Responds to Ping** - Select **Yes**.
 - **Trust Level** - Select **Trusted**.
12. Click **OK**.
13. Click **Send Changes**.

Step 6.5. 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 **Force**.

Do not use a **Failsafe** network activation when changing the management IP address.

Step 6.6 Switch the NextGen Admin endpoint to the secondary NextGen Firewall

To access the secondary NextGen Firewall F-Series you can either [configure a Client-to-Site VPN](#) on the primary NextGen Firewall F-Series to access the MIP of the secondary NextGen Firewall F-Series, or switch the MGMT (TCP port 807) Endpoint to the secondary firewall:

```
Get-AzureVM -ServiceName YOUR_CLOUD_SERVICE -Name  
YOUR_PRIMARY_BARRACUDA_NG_FIREWALL | Remove-AzureEndpoint -Name "MGMT" |  
Update-AzureVM Get-AzureVM -ServiceName YOUR_CLOUD_SERVICE -Name  
YOUR_SECONDARY_BARRACUDA_NG_FIREWALL | Add-AzureEndpoint -Protocol tcp -  
LocalPort 807 -PublicPort 807 -Name "MGMT" | Update-AzureVM
```

Step 6.7 Reconfigure the secondary firewall

Complete [Steps 6.1 - 6.5](#) for the secondary firewall.

If you are using PAYG (hourly) images you must export the license files of the secondary firewall, as only these licenses validate on the secondary firewall.

Both Barracuda NextGen Firewall F-Series systems are now using the static 'eth0' network interfaces (**CONTROL > Network**).

Interface/IP	Label	Ping	MAC of duplicate IP	Info
eth0				
10.0.20.6/24	net1	ok	-	
lo				

Table / Src Filter	State	Type	Interface	Src IP	Pref	Gateway	Name
Table main. From all							
10.0.20.0/24	up	direct-k...	eth0	10.0.20.6	0	-	IPAD01
127.0.0.0/24	up	direct-b...	lo	127.0.0.2	0	-	boxnet
Table default. From all							
0.0.0.0/0	up	gateway...	eth0	10.0.20.6	0	10.0.20.1	ROUT01

Step 7. Create a DHA cluster configuration

Create a DHA cluster configuration. Switch the Endpoint to the primary NextGen Firewall F-Series (see Step 6.6.).

1. Log into the primary Barracuda NextGen Firewall F-Series.
2. Go to **CONFIGURATION > Configuration Tree**.
3. Right-click on **Box** and select **Create DHA Box**.
4. Go to **CONFIGURATION > Configuration Tree > HA Box > HA Network**.
5. Select **eth0** from the **Interface Name** list.
6. Enter the static IP address of the secondary Barracuda NextGen Firewall F-Series as the **Management IP (MIP)**. E.g., 10.0.20.7
7. In the left menu, click **Routing**.
8. Verify the default route is present. (0.0.0.0/0 gateway XX.XX.XX.1).
9. In the left menu, click on **IP Configuration**.
10. In the **Additional Local IPs** section, click + to add a **IP Address Configuration**. The **IP Address Configuration** window opens.
11. Enter a **Name**
12. Click **OK**.
13. Enter the following settings for the **IP Address Configuration**:
 - o **Interface Name** - Select **eth1**.
 - o **IP Address** - Enter the IP address you assigned to this additional NIC of the secondary firewall during deployment.
 - o **Associated Netmask** - Select the netmask for this subnet.
 - o **Responds to Ping** - Select **Yes**.
 - o **Trust Level** - Select **Trusted**.
14. Click **OK**.
15. Click **Send Changes** and **Activate**.

Step 8. Set the primary and secondary firewall

Configure which NextGen Firewall F-Series is used as the active and backup box for the virtual server.

1. Log into the primary Barracuda NextGen Firewall F-Series.
2. Go to **CONFIGURATION > Configuration Tree > Virtual Servers > your virtual server > Server Properties**).
3. Click **Lock**.
4. From the **Active Box** drop-down select **This-Box**.
5. From the **Backup Box** drop-down select **Other-Box**.
6. Click **Send Changes** and **Activate**.

Step 9. Deploy the HA PAR file to the secondary firewall

Step 9.1 Create the PAR file for the HA firewall.

1. Log into the primary Barracuda NextGen Firewall F-Series firewall.
2. Go to **CONFIGURATION > Configuration Tree**.
3. Right-click on **Box** and select **CREATE PAR FILE for HA box**. You are prompted to save the boxha.par file.

Step 9.2 Deploy the PAR file on the secondary firewall

Switch the endpoint to the secondary NextGen Firewall F-Series (see Step 6).

1. Log into the secondary Barracuda NextGen Firewall F-Series firewall.
2. Go to **CONFIGURATION > Configuration Tree**.
3. Right-click on **Box** and select **Restore from PAR file**.
4. Choose the boxha.par file created in Step 9.1.
5. Click **Activate**.
6. Go to **CONTROL > Box**.
7. In the left navigation in the **Network** section, click on **Activate new network configuration**.
8. Click **Failsafe**.
9. In the left navigation in the **Operating System** section, click **Firmware Restart**.

Step 10. (optional) Remove the SETUP-MGMT-ACCESS access rule

This redirect access rule is no longer needed and can be deleted.

1. Log in to the primary Barracuda NextGen Firewall F-Series.
2. Go to **CONFIGURATION > Configuration Tree > Virtual Servers > S1 > Firewall > Forwarding Rules**.
3. Click **Lock**.
4. Right-click on **SETUP-MGMT-ACCESS** firewall rule and click **Delete**.
5. Click **Send Changes** and **Activate**.

You can now use the Barracuda NextGen Firewall F-Series HA cluster in the Microsoft Azure cloud.

Step 11. Configure Azure User Defined Routing

Azure User Defined Routing allows you to use the NextGen Firewall F-Series HA cluster in the frontend subnet as the default gateway for all your VMs running in the backend networks. You must enable IP forwarding for the NextGen Firewall F-Series VMS and create and apply an Azure routing table to the backend networks. Using a management Certificate and the Azure subscriber ID the NextGen Firewall F-Series VMs can change the Azure Routing Table on the fly when the virtual server fails over from one VM to the other.

Step 11.1 Configure User Defined Routes for your VNET

Create a User Defined routing table and enable IP Forwarding for the two NextGen Firewall F-Series VMs. Assign this user defined routing table to all subnets that use the NextGen Firewall F-Series HA cluster as the default gateway.

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

Step 11.2 Create the Azure Management certificate

For the NextGen Firewall F-Series to be able to connect to the Azure backend, you must create and upload a Management certificate.

1. Log in to the NextGen Firewall F-Series via ssh .
2. Create the certificate:
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout mycert.pem -out mycert.pem
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout mycert.pem -out mycert.pem
3. Answer the questions at the prompt. The **Common Name** is used to identify this certificate in the Azure web interface.
4. Convert the certificate to CER required by Azure:
openssl x509 -inform pem -in mycert.pem -outform der -out mycert.cer

```
openssl x509 -inform pem -in mycert.pem -outform der -out mycert.cer
```

You now have two certificates mycert.pem and mycert.cer.

Step 11.3 Upload the Azure Management certificate

1. Log into the Microsoft Azure Management Portal (<https://manage.windowsazure.com>).
2. On the bottom of the left menu, click on **SETTINGS**.
3. In the top navigation, click on **MANAGEMENT CERTIFICATES**.
4. On the bottom Click **UPLOAD**.
5. Select the *mycert.cer* certificate created in Step 12.2. and click **OK**.

The management certificate is now listed with the **Common Name** of the certificate used as the **Name**.

Step 11.4. Configure Cloud Integration

You must enter your Azure SubscriptionId, VNET name and the management certificate to allow the F-Series Firewall to change the Azure User Defined Routing Table.

1. Login to the primary firewall.
2. Go to **CONFIGURATION > Configuration Tree > Box > Advanced Configuration > Cloud Integration**.
3. Click **Lock**.
4. In the left menu, click **Azure Networking**.
5. Enter your Azure **Subscription ID**. Use `Get-AzureSubscription` in Azure PowerShell to display your SubscriptionId.
6. Enter the **Virtual Network Name**.
7. Next to **Management Certificate** click **Ex/Import** and select **Import from PEM File**. The **File browser** window opens.
8. Select the *mycert.pem* certificate created in step 12.2 and click **Open**.
9. Next to **Management Key** click **Ex/Import** and select **Import from File**. The **File browser** window opens.
Select the *mycert.pem* certificate created in step 12.2 and click **Open**.
If you are using an OpenSSL version that generates PKCS#8 keys import the *mycert.key.pem* file as the **Management Key** on the F-Series Firewall.
10. Set **Protect IP forwarding settings** to **yes**.

Azure Networking	
Azure Deployment Type	<input type="text" value="Azure-Service-Management-(ASM)"/>
Subscription ID	<input type="text" value="bde58b49-9951-466e-90e2-592c0920ce77"/>
Tenant ID	<input type="text"/>
Application ID	<input type="text"/>
Resource Group	<input type="text"/>
Virtual Network Name	<input type="text" value="wideVNET"/>
Route Check Interval	<input type="text" value="300"/>
Management Certificate	<input type="button" value="Show..."/> <input type="button" value="Ex/Import"/> Hash: QOEUXZ 1024 Bits
Management Key	<input type="button" value="New Key..."/> <input type="button" value="Ex/Import"/> Hash: NOSMDC 2048 Bits
Protect IP forwarding settings	<input type="text" value="yes"/>

11. Click **Send Changes** and **Activate**.

The Azure Routing table is now updated every time the virtual server fails over.

Next steps

- Activate the license on both Barracuda NextGen Firewalls. For more information, see [How to Activate and License a Stand-alone Virtual or Public Cloud F-Series Firewall or Control Center](#).
- Enable Azure Route Table rewriting. For more information, see [How to Configure Azure Cloud Integration using ASM](#).

Figures

1. ha_setup_azure.png
2. AzureMN_01.png
3. AzureMN_02.png
4. AzureHA07.png
5. AzureHA08_V2.png
6. Azure_default_route2.png
7. AzureHA11.png
8. UDR_HA_01.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.