How to Configure a High Availability Cluster in Azure with the Standard Load Balancer

https://campus.barracuda.com/doc/84313265/

Configure a high availability cluster to ensure that the services running on the Barracuda CloudGen Firewall 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 either an Availability Set or Availability Zone (where available). This ensures that the VMs are placed in different fault and update domains inside the Azure data center. 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 fail over immediately after the service 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 standard load balancer allows stateful sessions to remain as there are no IP address changes with this method. 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, a Standard type internal load balancer IP address is used, and this load balancer directs traffic to the active firewall. Now, the backend VMs can connect via the active firewall to the Internet.

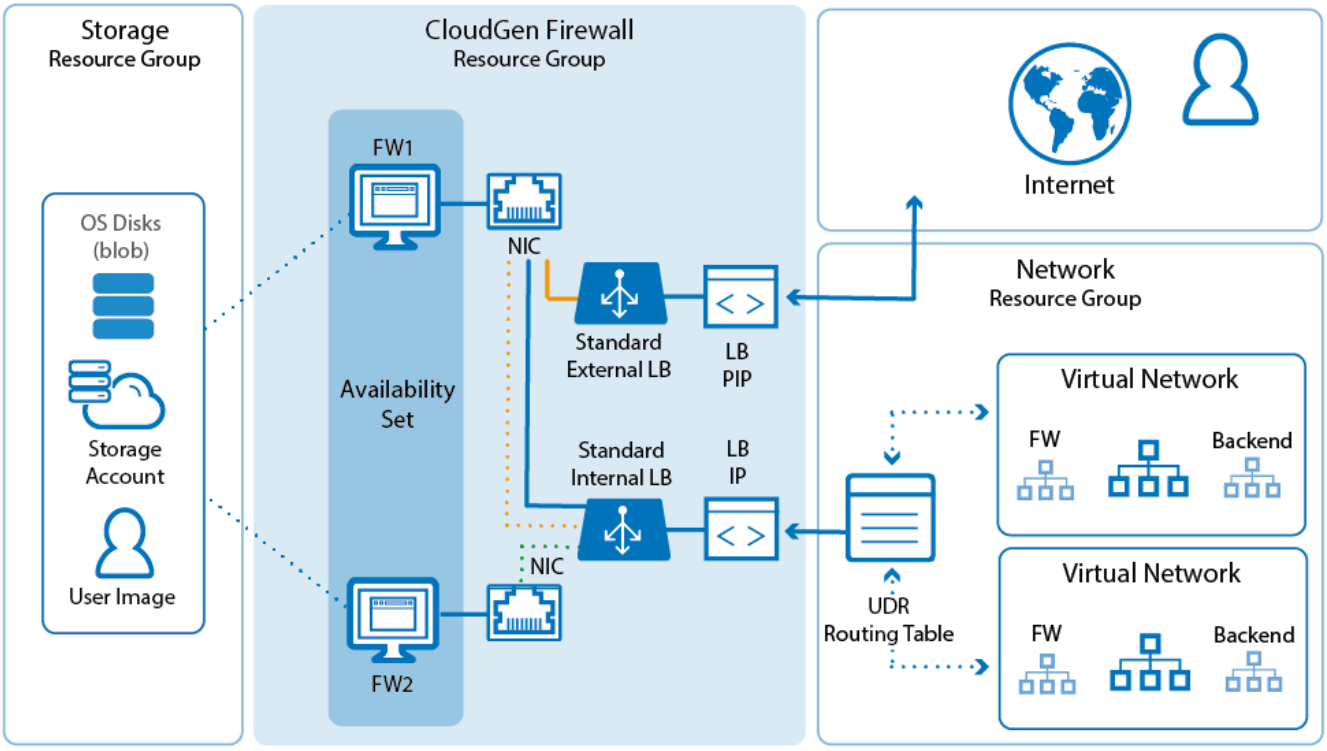

Step 1. Deploy Two CloudGen Firewall VMs

To configure an HA cluster, deploy two CloudGen 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:

- Static private (internal) IP addresses must be used.
- The SKU of the public IP of each firewall must match the SKU of the load balancer. In this case, the public IP must have a standard SKU since a standard load balancer will be used.
- The same instance size for both VMs must be used.
- Both firewalls must be the same Barracuda CloudGen Firewall for Azure model.
- Both VMs must be deployed in one Availability Set or across Availability Zones.

For more information, see How to Deploy a CloudGen Firewall from the Microsoft Azure Marketplace or How to Deploy a CloudGen Firewall in Microsoft Azure Using PowerShell and ARM.

Official templates are available to assist you to deploy quicker. These can be found in our GitHub: https://github.com/barracudanetworks/ngf-azure-templates . If you are using the GitHub template, provisioning may take a while. Until it completes, you will get the error message "access denied" if you try to connect via Barracuda Firewall Admin. If boot diagnostics are enabled, you can view the log. Further deployment examples can be found in the *contrib* folder.

To test, you can deploy a stand-alone proof of concept environment by following this guide: https://www.barracuda.com/resource/ref_architectures/azure_high_availability_cluster.

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.

> You can skip this when deploying clustered templates as these pre-complete these steps.

1. Go to **CONFIGURATION > Configuration Tree > Box > Network**.

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

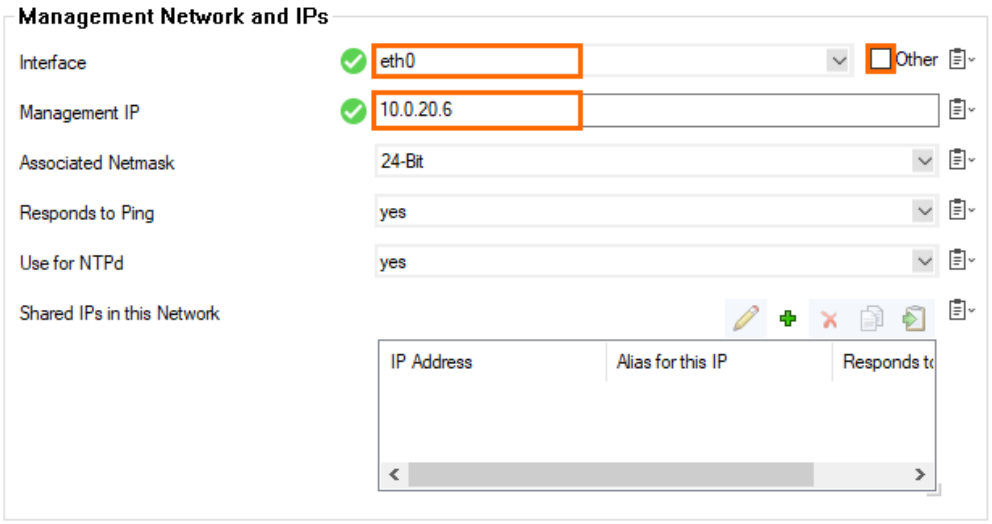

Step 2.2 Create the Default Route

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.

> You can skip this when deploying clustered templates as these pre-complete these steps.

1. In the left menu, click **Advanced Routing**.
2. Click **+** in the **IPv4 Routing 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**.

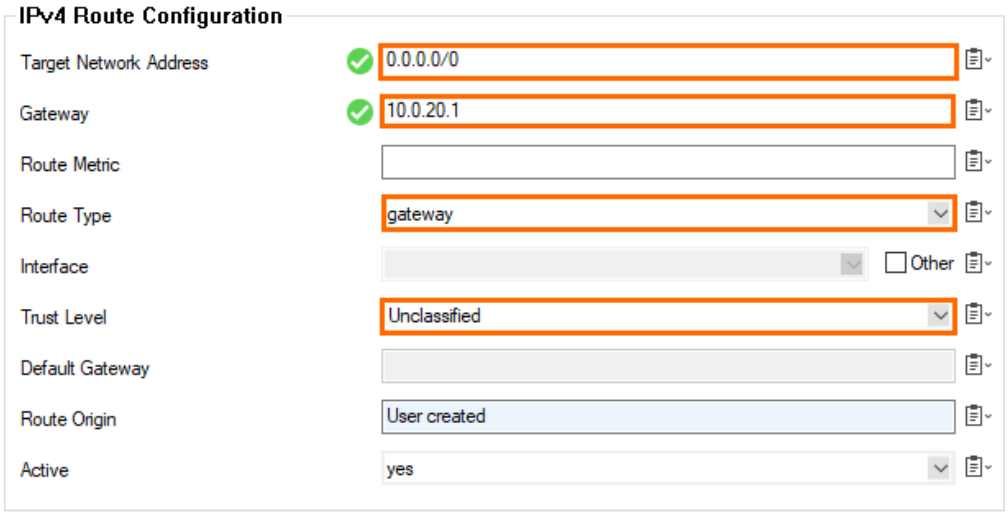

3. Click **OK**.
4. Click **Send Changes** and **Activate**.

Step 2.3 Disable ICMP Monitoring of the Gateway

ICMP probing must be disabled for the interface.

1. Go to **CONFIGURATION > Configuration Tree > Infrastructure Services > Control**.
2. Click **Lock**.
3. In the **ICMP Gateway Monitoring Parameter** section, click **+** to add an entry to the **No Probing for Interface** table.

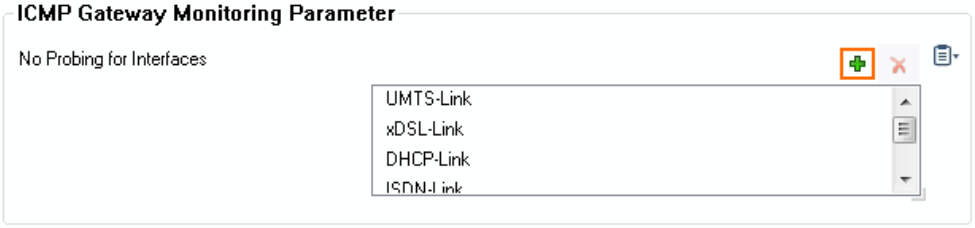

4. In the **Other** field, enter `eth0` .

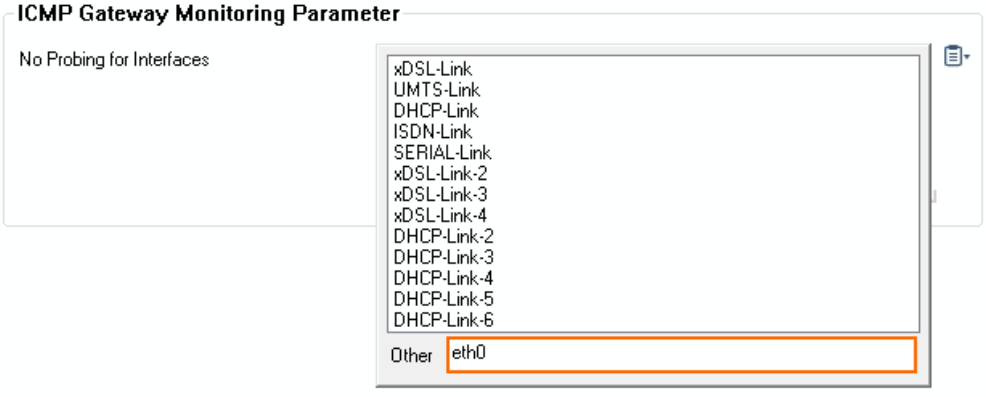

5. Click **Send Changes** and **Activate**.

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 **Activate new network configuration**.
3. Click **Failsafe**.

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

Step 3. (PAYG only) Import PAYG Licenses from the Secondary Firewall

Step 3.1 Export the PAYG license from the Secondary Firewall

1. Log into the secondary firewall.
2. Go to **CONFIGURATION > Configuration Tree > Box > Box Licenses**.
3. Click **Lock**.
4. Select the license file, click **Export**, and select **Export to File**.
5. Click **Unlock**.

Step 3.2 Import the PAYG License on the Primary Firewall

1. Log into the primary firewall.
2. Go to **CONFIGURATION > Configuration Tree > Box > Box Licenses**.
3. Click **Lock**.
4. Click **+** and select **Import from Files...**.
5. Select the license file exported from the secondary firewall.
6. Click **OK**.
7. Select **I agree** to accept the terms and conditions.
8. Click **OK**.
9. Click **Send Changes** and **Activate**.

The primary firewall now has both PAYG licenses listed in the **Licenses** list.

Step 4. Configure an HA Cluster on the CloudGen 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 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 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 service is

currently running on.

For more information, see How to Set Up a High Availability Cluster.

Step 5. (BYOL only) Activate and License the two Firewall VMs

Activate the license on the secondary firewall, then on the primary firewall. If the primary unit is activated prior to the secondary unit, the licenses for the secondary cannot be downloaded. In this case, reboot the primary firewall, perform a complete manual HA sync, and update to download and install the licenses correctly.

For more information, see How to Activate and License a Standalone High Availability Cluster.

Step 6. Create the External Load Balancer

After you have deployed the two firewalls, you can create the load balancers that will direct traffic as required. You will need to create a new external and internal load balancer to handle all potential traffic flows.

From the Azure portal:

1. Go to the upper left-hand corner and click the **+** symbol.
2. Type in `Load Balancer` and press **Enter**. The **Create Load Balancer** page opens.
3. Click **Create**.

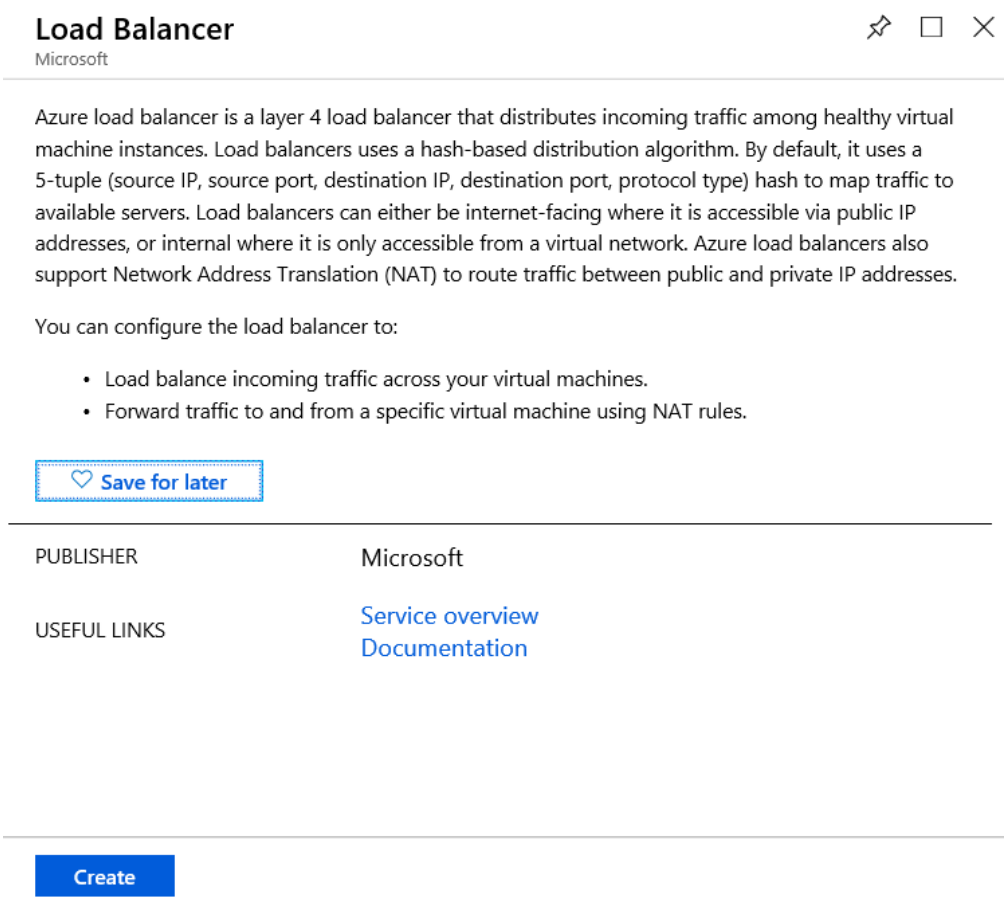

4. On the next page, configure the following settings:
   - **Name** – Enter the name of the External Load Balancer.
   - **Type** – Select the **Public** type.
   - **SKU** – Select the Standard SKU.
   - **Public IP address** – Create a new one for use with the load balancer and enter a name. Its SKU must be standard.
   - **Availability zone** – Select **Zone-redundant**.

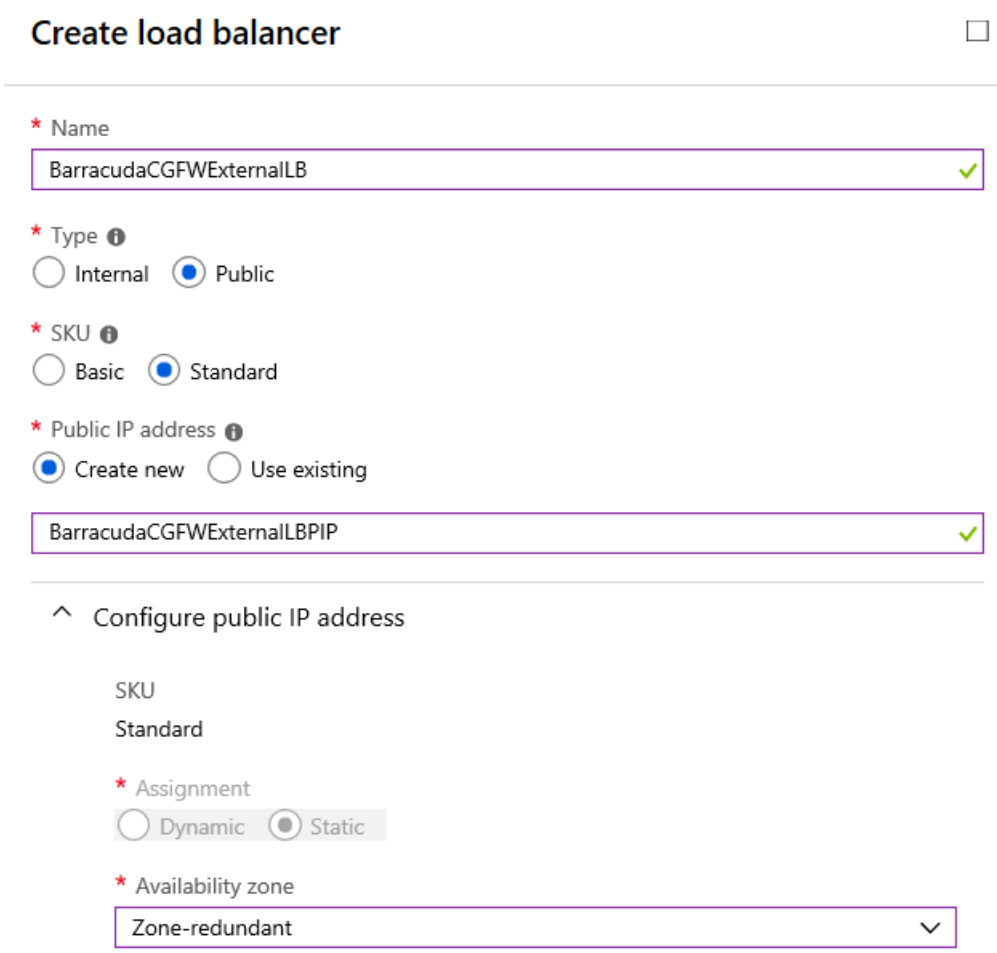

- **Subscription** – Select the Azure subscription.
- **Resource group** – Enter a unique name for your resource group, or click **Use Existing** and select an existing resource group.
- **Location** – Select the location of the firewall.

5. Navigate to your newly created load balancer.
6. Within the load balancer, go to **Health Probes** and click to **Add** a new one.
   - **Name** – As desired, but leave the probe settings as pictured below.
   - **Protocol** – **TCP**
   - **Port** – `65000`
   - **Interval** – Leave as default.
   - **Unhealthy threshold** – Leave as default.

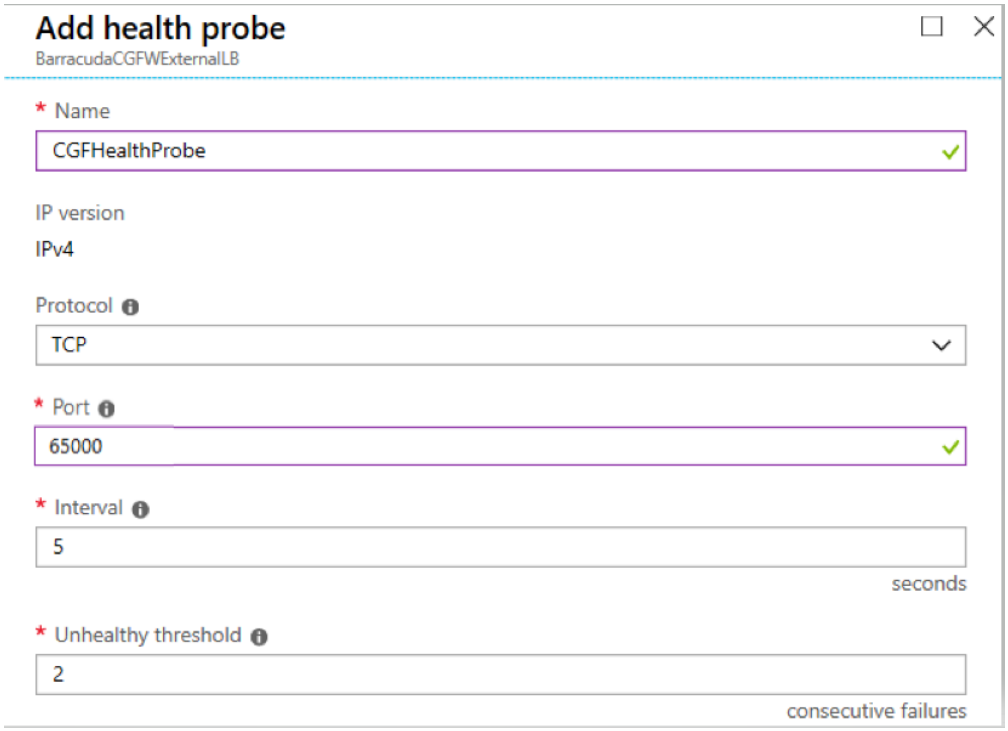

7. Click **OK** to create the probe.
8. Back in the Load Balancer, go to **Backend Pools** and click to **Add** a new one. Complete the fields as below:
   - **Name** – Enter your desired name for the Backend Pool.
   - **Virtual Network** – Select the Virtual Network you built your Firewalls in.
   - **Virtual Machine** – Select the first firewall you built.
   - **IP Address** – Select the ipconfig for the IP you wish the LB to send traffic to.
   - Repeat for the second Firewall VM you built.
9. Click **Add**.
10. To create the load balancing rules for inbound traffic, you must have one for TCP and one for UDP in order for the firewalls to pass traffic out to the Internet on those ports. This instruction creates rules for inbound client VPNs that meet this requirement. Go to **Load Balancing Rules** and click **Add** to create a new rule.
    - **Name** – Suggested name `TINA-TCP`
    - **IP Version** – **IPv4**
    - **Frontend IP address** – Select the front-end IP created at build time.
    - **Protocol** – **TCP**
    - **Port** – 691
    - **Backend port** – 691
    - **Backend pool** – Select the backend pool just created.
    - **Health probe** – Select the probe created previously.
    - **Session persistence** – Leave as default.
    - **Idle timeout (minutes)** – Leave as default.
    - **Floating IP (direct server return)** – Leave as default.

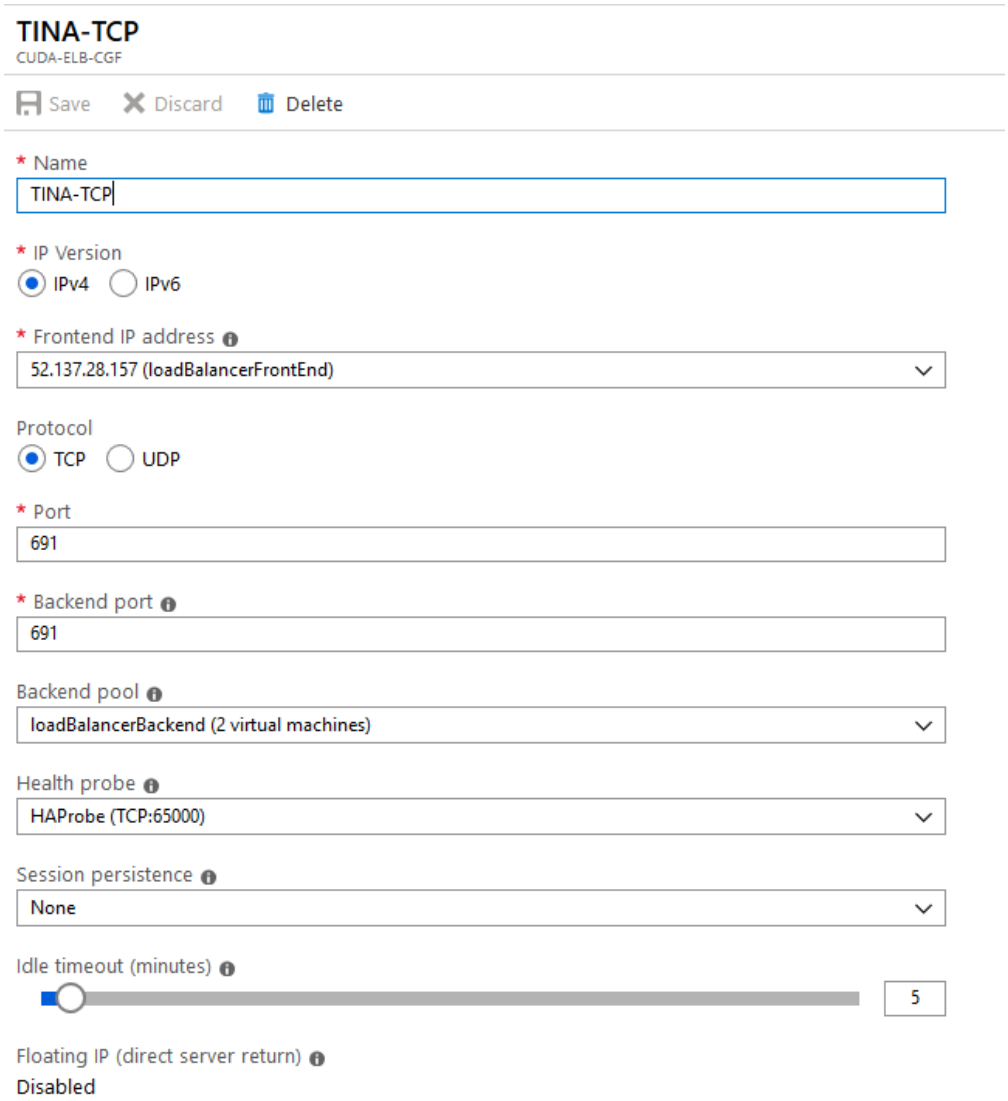

11. Click **OK** to create the rule.
12. Repeat the steps above to create a second rule for UDP, but change the following settings:
    - **Name** – Suggested name TINA-UDP
    - **Protocol** – **UDP**
    - **Port** – 691
    - **Backend port** – 691

Step 7. Create the Internal Load Balancer

The internal load balancer is essential for a standard load balancer HA design because it is the destination for all user-defined routes. To create one:

From the Azure portal:

1. Go to the upper left-hand corner and click the **+** symbol.

2. Type in `Load Balancer` and select the option labeled **Load Balancer**.
3. Click **Create**.

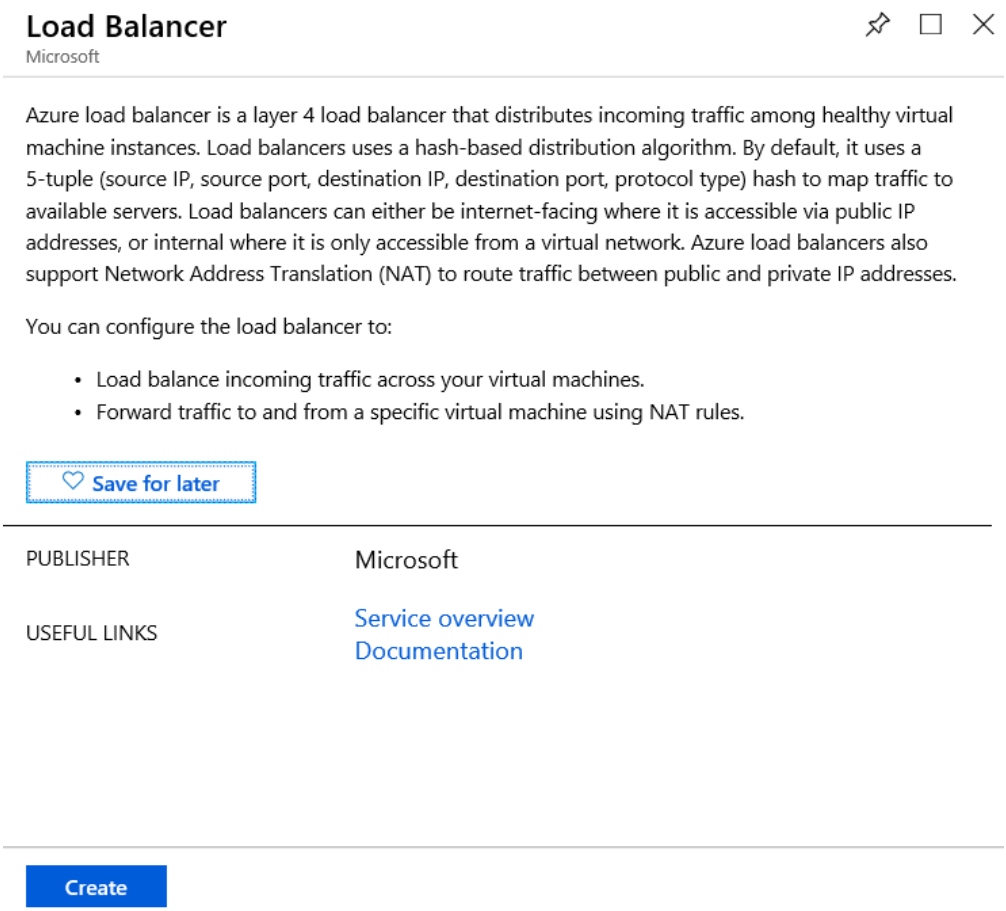

4. On the next page, configure the following settings:
   - **Name** – Enter the name of the internal load balancer.
   - **Type** – Select the **Internal** type.
   - **SKU** – Select the Standard SKU.
   - **Virtual Network** – Select the virtual network your firewalls are in.
   - **Subnet** – Select the subnet the firewalls are in.
   - **IP Address assignment** – **Static**

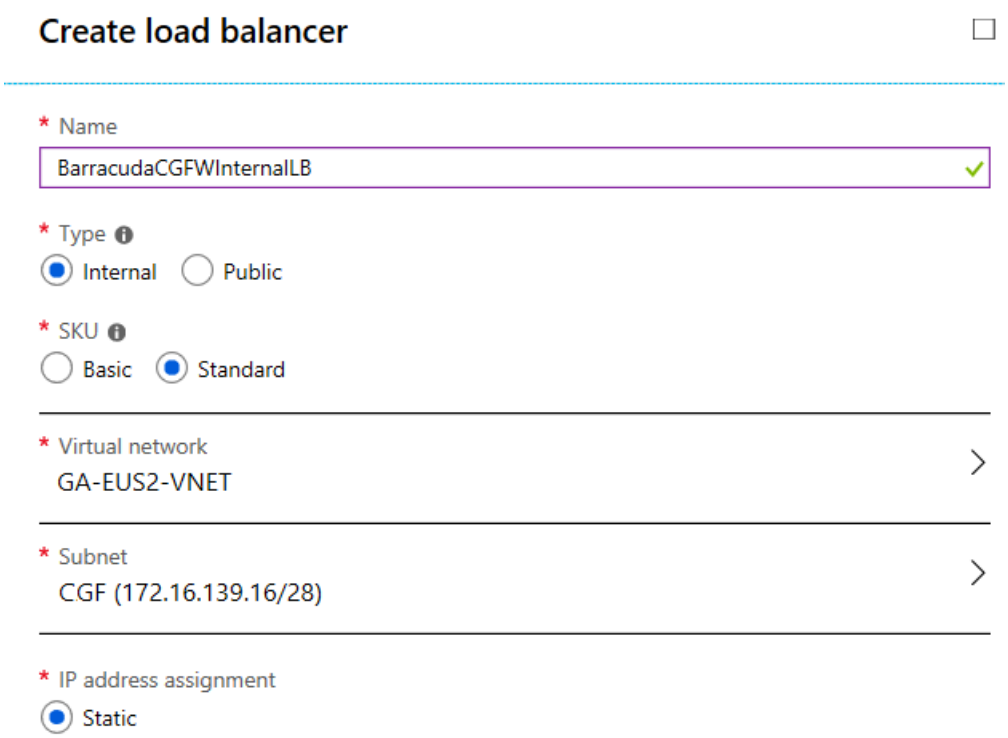

- **Private IP address** – Enter a private IP in that subnet for the load balancer to use.
- **Availability Zone** – Select **Zone-redundant**.
- **Subscription** – Select the Azure subscription.
- **Resource group** – Enter a unique name for your resource group, or click **Use Existing** and select an existing resource group.
- **Location** – Select the location of the firewall.

5. Navigate to your newly created load balancer.
6. Within the load balancer, go to **Health Probes** and click to **Add** a new one.
   - **Name** – As desired, but leave the probe settings as pictured below.
   - **Protocol** – **TCP**
   - **Port** – `65000`
   - **Interval** – Leave as default.
   - **Unhealthy threshold** – Leave as default.

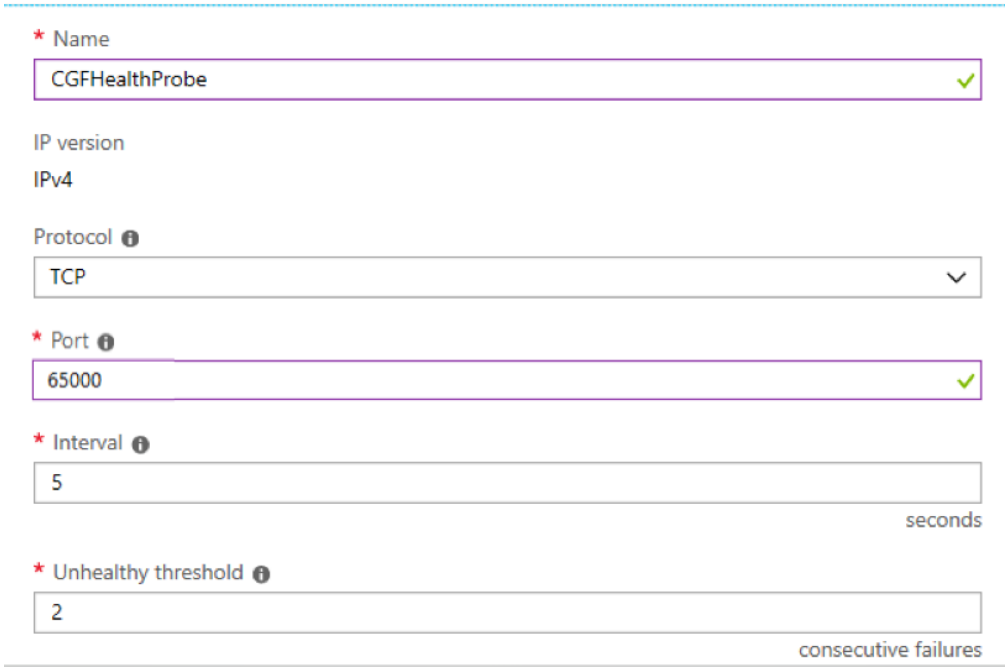

7. Click **OK** to create the probe.
8. Back in the load balancer, go to **Backend Pools** and click to **Add** a new one. Complete the fields as below:
   - **Name** – Enter your desired name for the backend pool.
   - **Virtual Network** – Select the virtual network you built your firewalls in.
   - **Virtual Machine** – Select the first firewall you built.
   - **IP Address** – Select the ipconfig for the IP you wish the load balancer to send traffic to.
   - Repeat for the second firewall VM you built.
9. Click **Add**.
10. Create the load balancing rules for traffic to flow. (You do not need to define ports with this type of internal load balancer.)
    - **Name** – Suggested name `AllPortsHA`
    - **IP Version** – **IPv4**
    - **Frontend IP Address** – This will be the private IP allocated on creation.

      

    - **HA Ports** – Select.
    - **Backend Pool** – Select the backend pool just created.
    - **Health Probe** – Select the probe created previously .
    - **Session Persistence** – Leave as default.
    - **Idle Timeout** – Leave as default.
    - **Floating IP (direct server return)** – Leave as default.
11. Click **OK** to create the rule.
12. Now you have completed the setup of the load balancers.

Step 8. Enable IP Forwarding

To allow the firewall to pass traffic not intended for itself, you must update the network interface.

In the Azure portal,

1. Go to the virtual machine.
2. Go to **Networking**, and locate the Network Interface attached to the firewall.
3. In **IP Configurations**, make sure that **IP Forwarding** is enabled.

If not already done, make the ipconfig static by clicking on it and setting the assignment to **static**.

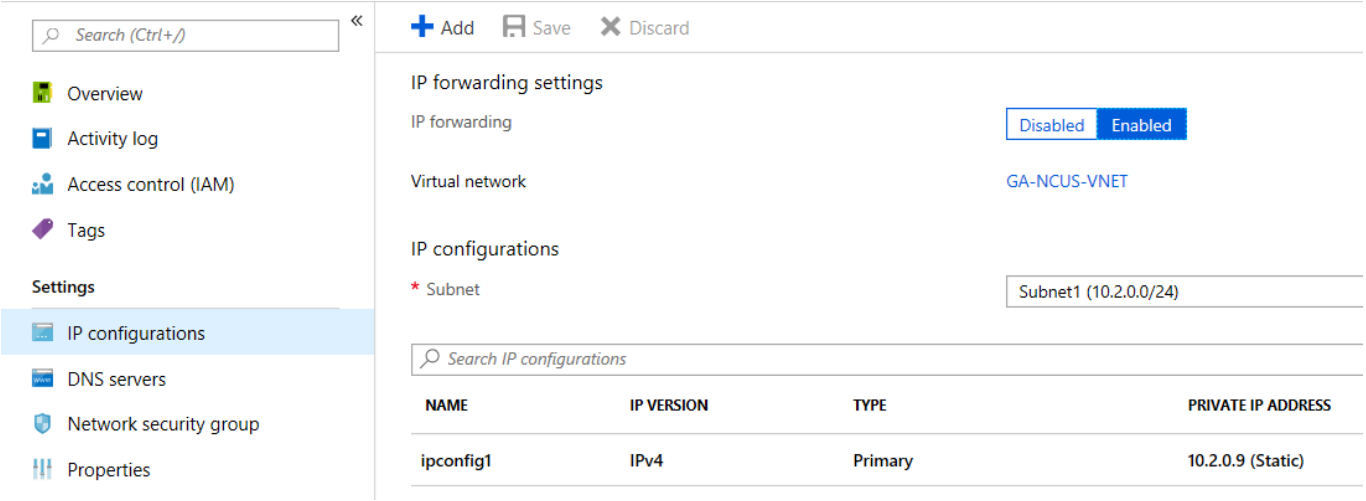

For more information, see How to Configure a Client-to-Site VPN Group Policy or How to Configure a Client-to-Site TINA VPN with Personal Licenses .

Step 9. Allow the Load Balancer Health Probes to Succeed

Activate the preconfigured firewall forwarding rule to allow load balancer health probes to succeed. The connection will use the port you indicated in Steps 2 & 3 above. It will originate from 168.63.129.16 and can be redirected to any service running locally on the firewall (e.g., 127.0.0.9:450 for firewall authentication service, or 127.0.0.9:691 for FW TINA VPN).

1. Log into the Barracuda CloudGen Firewall with Barracuda Firewall Admin.
2. Go to **CONFIGURATION > Configuration Tree > Assigned Services > NGFW > Forwarding Rules**.
3. Click **Lock**.

4. Open the rule CLOUD-LB-PROBE.
5. To activate the rule, clear the check box next to **Deactivate Rule**.

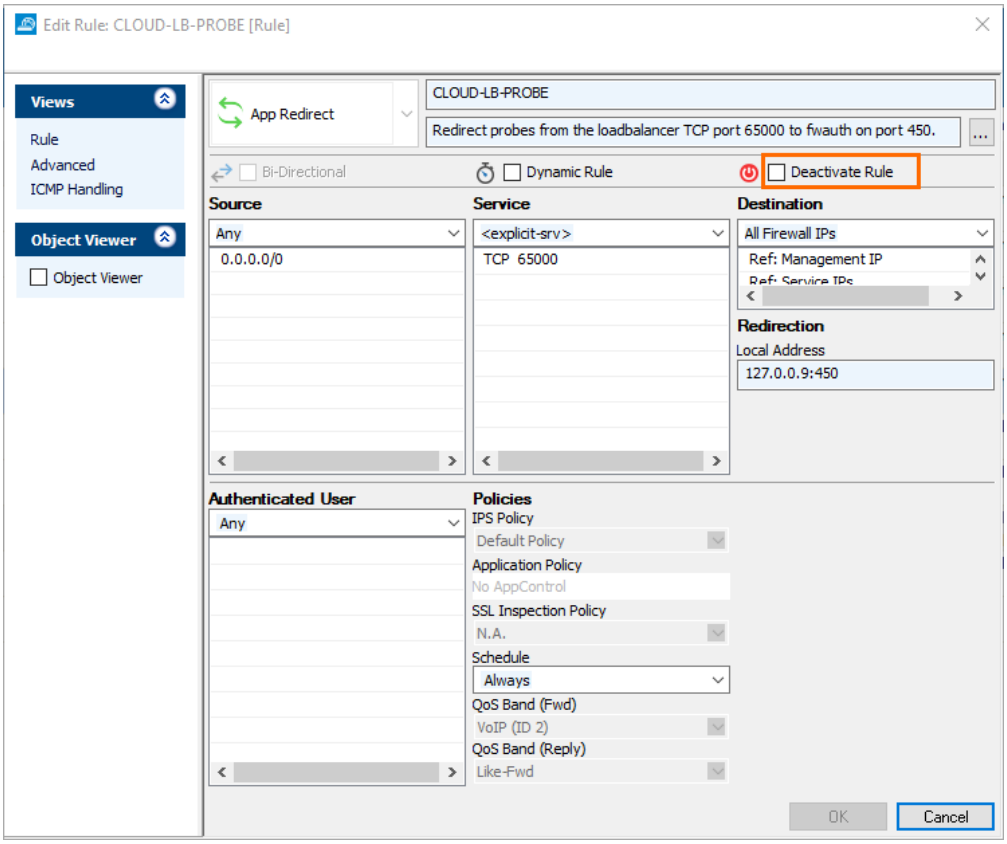

6. Click **OK**.
7. Click **Send Changes** and **Activate**, then click **Activate**.

Step 10. Configure User-Defined Routing in Azure

Configure UDR for the backend VMs to use the internal load balancer's IP as their default gateways for all connections to the Internet. 0.0.0.0/0 will only impact traffic that does not have a route already present in Azure. E.g., Internet.

To affect traffic within the VNET, subnet, or peered VNET, introduce routes for a matching destination network. (Check the effective routing of a VM if uncertain what routes are present already).

Step 11. 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 VPN Group Policy or How to Configure a Client-to-Site TINA VPN with Personal Licenses .

Step 12. Disassociate the Public IP Addresses

When both a load balancer 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 means that outgoing connections use different source IP addresses depending on which firewall is active.

> If the portal deployed the box using a basic public IP on the network interface, you must remove it for the external standard load balancer to work correctly. You can replace it with standard SKU public IPs directly on the NIC.

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 **Public IP Address**. The **Public IP address** column opens.
4. Click **Disassociate**.

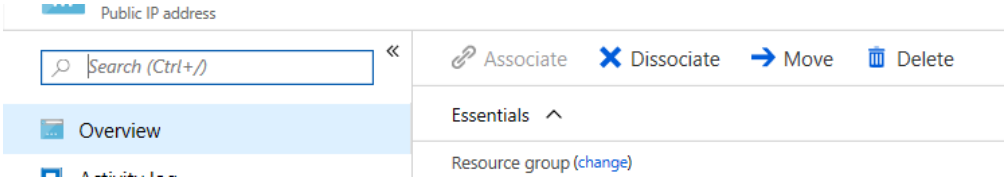

5. Repeat for the secondary firewall VM.

Use a client-to-site VPN connection to manage both Barracuda CloudGen Firewall VMs via the internal IP addresses. For more information, see Client-to-Site VPN.
Go to the **Firewall > History** view and confirm you can see the health probes succeeding. Traffic should be passing through the firewall correctly. If you see timeouts, confirm NSGs on the interfaces permit traffic and that **IP Forwarding** is enabled.

Example of Successful Monitoring Polls on port 65000 or 691

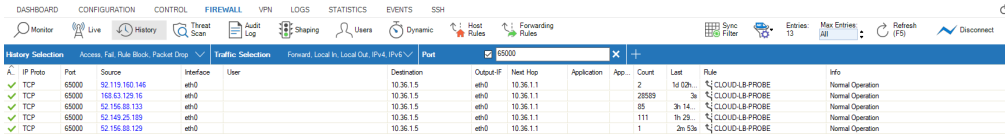

Figures

1. azure_std_ha_diagram.png
2. MIP.png
3. routeipv4.png
4. disable_icmp_probing_01.png
5. disable_icmp_probing_02.png
6. lbs01.png
7. lbs02.png
8. lb_health_probe_port65k.png
9. lb_rule.png
10. lbs05.png
11. create LB_cgf.png
12. lb_health_probe_port65000.png
13. lbs08.png
14. lbs11.png
15. lb_probe_rule_acitvate.png
16. lbs12.png
17. health_probe_history_view.png