

# Enterprise Azure Configuration Guide

Version 8.13.2 Revision 1.0.0



# Table of Contents

1. Introduction . . . . .	3
2. About Enterprise Azure . . . . .	3
2.1. Main Differences to Our Standard (Non-Cloud) Product . . . . .	3
2.2. Why use Enterprise Azure? . . . . .	4
3. Azure Deployment Models . . . . .	4
4. Accessing Microsoft Azure . . . . .	5
5. Azure Management . . . . .	5
5.1. Accessing the Azure Portal . . . . .	5
5.2. Azure PowerShell & Azure CLI . . . . .	5
6. Deploying Enterprise Azure From the Marketplace . . . . .	5
7. Accessing the Appliance . . . . .	11
7.1. Accessing the Appliance WebUI . . . . .	11
7.1.1. WebUI Menu Options . . . . .	13
7.2. Appliance Security . . . . .	13
7.2.1. Security Mode . . . . .	14
7.2.2. Passwords . . . . .	14
7.3. Appliance Software Update . . . . .	14
7.4. Appliance Licensing . . . . .	15
7.5. Enterprise Azure Non-standard WebUI Menu Options . . . . .	15
7.6. Accessing the Appliance using SSH . . . . .	16
7.6.1. Accessing the Appliance from Linux . . . . .	16
7.6.2. Accessing the Appliance from Windows using PuTTY . . . . .	16
8. High Availability in Azure . . . . .	17
8.1. Key Concepts . . . . .	17
8.2. Implementing HA in Azure . . . . .	18
9. Configuration Examples . . . . .	19
9.1. Example 1 - Load Balancing Web Servers, HA Configuration, 1 Subnet, Layer 7 . . . . .	19
9.2. Example 2 - Load Balancing Web Servers, Single Appliance, 2 Subnets, Layer 4 NAT Mode . . . . .	29
10. Testing & Verification . . . . .	35
11. More Information . . . . .	35
12. Loadbalancer.org Technical Support . . . . .	35
12.1. Contacting Support . . . . .	35

# 1. Introduction

Microsoft Azure is a comprehensive set of cloud services that developers and IT professionals use to build, deploy and manage applications through Microsoft's global network of data centers. It allows services to be deployed as and when required. Charges are made for what is used making it an extremely flexible and cost effective solution. The Loadbalancer.org Enterprise Azure cloud based load balancer allows customers to rapidly deploy and configure a feature rich load balancing solution within the Azure cloud.

## 2. About Enterprise Azure

Enterprise Azure is a fully featured Application Delivery Controller (ADC) / load balancer designed specifically for Azure. The core software is based on LBOS-7 which is a customized Linux build maintained by Loadbalancer.org, LVS, Ldirectord, Linux-HA, HAProxy & STunnel.

Enterprise Azure can be deployed as a single instance although we always recommend that 2 appliances are deployed as an HA clustered pair to avoid introducing a single point of failure.

Enterprise Azure is based on the same code base as our main hardware/virtual product. This means that Enterprise Azure supports many of the same features as the hardware & virtual based products. There are certain differences due to the way the Microsoft Azure environment works. The main differences are listed below.

### 2.1. Main Differences to Our Standard (Non-Cloud) Product

1. Layer 4 DR mode is **not** currently supported.
2. In Azure, you should configure your HA pair **first** before setting up your load balanced services. This is different to the recommendation for our hardware/virtual products and is due to the way HA is handled. In our standard product, when a failover occurs, the **same** VIP address is brought up on the passive device. In Azure, in order to minimize the time taken for the failover a different approach is used. When creating a VIP on an Azure HA pair, 2 private IPs must be specified - one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. The IPs for the VIP on the Primary & Secondary are selected using drop-downs within the VIP configuration screen. An Azure load balancer is used in front of the Loadbalancer.org HA pair to direct inbound traffic to the active appliance. Both Primary & Secondary appliances must be in the same Availability Set or deployed within the same Availability Zone or split across 2 different Availability Zones. Please refer to [High Availability in Azure](#) for more information on configuring an HA pair.

The private IPs for the VIP on the Primary & Secondary are selected using drop-downs within the VIP configuration screen. These drop-downs are only displayed **after** the pair is configured. They are populated with the IPs that are assigned to the network interface using the WebUI option: *Local Configuration > Network Interface Configuration*.

#### Note

**Adding VIPs after creating an HA pair (RECOMMENDED)** - If you add VIPs after creating an HA pair, you'll be prompted for both IPs. Add the IPs you intend to use for the VIPs to the local interface on both Primary & Secondary and they'll be available in the drop-downs.

**Creating an HA pair after configuring VIPs on the Primary** - If you add a Secondary appliance and create an HA pair after adding VIPs to the Primary appliance, the floating IPs that were automatically configured for each VIP must first be removed using the WebUI



option: *Cluster Configuration > Floating IPs* and then added to the network interface instead. This will ensure that these IPs appear in the drop-downs mentioned above. You'll also need to configure IPs in a similar way on the Secondary device so that corresponding Secondary IPs can be selected for each VIP using the drop-downs.

3. Layer 4 NAT mode where the default gateway on the load balanced real servers is required to be the load balancer is **not** supported. Routing rules for the real server subnet must be changed instead - please refer to [Configuration Example 2](#) for more details.
4. Layer 7 SNAT mode with TProxy enabled where the default gateway on the load balanced real servers is required to be the load balancer is **not** supported. Routing rules for the real server subnet must be changed instead - please refer to [Configuration Example 2](#) for more details.
5. Layer 4 NAT mode and layer 7 SNAT mode with Transparent Proxy are not supported when using an HA clustered pair. In both cases, the custom routing rules would need to be dynamically modified to route via the Secondary appliance rather than the Primary if a failover occurs. This is currently not supported.

## 2.2. Why use Enterprise Azure?

- **Comprehensive features** - Enterprise Azure supports a wide range of features:
  - Supports comprehensive Layer 7 load balancing.
  - Load balances both Azure based and non-Azure based servers.
  - Supports Round Robin and Least Connection connection distribution algorithms.
  - Supports customizable timeouts for custom applications beyond those offered by Azure.
  - Supports comprehensive back-end server health-check options.
  - Enables fallback servers to be configured and invoked when all load balanced servers/services fail.
  - Provides extensive real time and historical statistics reports.
  - Supports session distribution based on actual server load (utilizing Loadbalancer.org's feedback agent which is available for both Linux & Windows).
  - Supports SSL Termination.
  - Supports Microsoft RDP Cookie based persistence.
  - Supports full integration with Microsoft Remote Desktop Services Connection Broker.
  - GSLB for multisite load balancing.
  - Fully featured WAF (Web Application Firewall).
- **Ease of use** - The interface is virtually identical to our hardware/virtual product which is very simple and intuitive to use. This also makes migrations to Azure much easier for existing customers.
- **Freedom license** - Our freedom license enables customers to migrate from one environment to another (e.g. virtual to cloud) at no additional cost - and with free migration assistance.
- **Expert assistance is available 24 x 7** - Our highly experienced support team can assist when needed.

## 3. Azure Deployment Models



The Azure platform currently supports both the original Classic model and the latest Resource Manager model. To simplify the deployment and management of resources, Microsoft recommends that the Resource Manager model is used for new resources, and, if possible, existing resources are re-deployed through Resource Manager. For a more detailed comparison of Classic and Resource Manager models, please click [here](#).

## 4. Accessing Microsoft Azure

To start using Microsoft Azure, you'll need an Azure account. If you don't already have one you can create one at the following URL: <https://azure.microsoft.com/en-us/get-started/>.

## 5. Azure Management

Azure resources can be managed in 3 ways:

- Azure Portal
- Azure PowerShell
- Azure CLI

### 5.1. Accessing the Azure Portal

The Azure Portal is available [here](#).

### 5.2. Azure PowerShell & Azure CLI

Information on how to obtain, install and configure Azure PowerShell is available [here](#).

Information on how to obtain, install and configure Azure CLI is available [here](#).

## 6. Deploying Enterprise Azure From the Marketplace

1. Login to the Azure Portal.
2. Access the **Marketplace** and search for "Loadbalancer.org", you'll be presented with the following options:
  - \*Loadbalancer Enterprise ADC - BYOL
  - \*Loadbalancer Enterprise ADC

 **Note**

The BYOL version will work completely unrestricted for 30 days without any license applied. During this period, only Azure Compute usage charges will apply. After the 30 days, the trial will still function, but no configuration changes will be possible until the license is applied.

3. Click on the option you require, you will be presented with a more detailed overview of the product.
4. Click the **Create** button.

### Configure the Basics



## Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription \* ⓘ

Loadbalancer.org Pay-As-You-Go

Resource group \* ⓘ

QS-RG1

[Create new](#)

## Instance details

Virtual machine name \* ⓘ

LB1

Region \* ⓘ

(Europe) UK South

Availability options ⓘ

Availability zone

Zone options ⓘ

Self-selected zone

Choose up to 3 availability zones, one VM per zone

Azure-selected zone (Preview)

Let Azure assign the best zone for your needs

Availability zone \* ⓘ

Zone 1

 You can now select multiple zones. Selecting multiple zones will create one VM per zone. [Learn more](#)

Security type ⓘ

Standard

Image \* ⓘ

 Load Balancer Enterprise ADC BYOL - x64 Gen1

[See all images](#) | Configure VM generation

VM architecture ⓘ

Arm64

x64

 Arm64 is not supported with the selected image.

Run with Azure Spot discount ⓘ



Size \* ⓘ

Standard\_F2s\_v2 - 2 vcpus, 4 GiB memory (£59.90/month)

[See all sizes](#)

Enable Hibernation ⓘ



 Hibernation is not supported by the image and size that you have selected. Choose an image and size that is compatible with Hibernation to enable this feature. [Learn more](#)



## Administrator account

Authentication type 

SSH public key  
 Password

 Azure now automatically generates an SSH key pair for you and allows you to store it for future use. It is a fast, simple, and secure way to connect to your virtual machine.

Username \* 

azureuser 

SSH public key source

Generate new key pair 

SSH Key Type

RSA SSH Format  
 Ed25519 SSH Format  
 Ed25519 provides a fixed security level of no more than 128 bits for 256-bit key, while RSA could offer better security with keys longer than 3072 bits.

Key pair name \* 

QS-LB1-keypair 

1. Configure the **Subscription & Resource group** settings according to your requirements.
2. Enter a suitable **Virtual machine name** for the instance, e.g. **LB1**.
3. Select the required **Region**.
4. Configure the **Availability options** according to your requirements.

 **Note**

For an HA clustered pair, both VMs must be in the same Availability Set or deployed within the same Availability Zone or split across 2 different Availability Zones. Please refer to [High Availability in Azure](#) for more details on setting up an HA pair.

 **Note**

For more details on Azure availability options, please refer to [Availability options for Azure Virtual Machines](#).

5. Leave the **Security Type** set to **Standard**.

6. Select the required **Size** - this can be changed by expanding the drop-down and selecting from the recently used or recommended image sizes. Or alternatively by clicking **See all sizes** and choosing from the expanded list of options.

 **Note**

The Image size required depends on the anticipated workload. For production deployments we recommend at least 2 vCPUs and 4GB RAM. For further help and advice please contact [support@loadbalancer.org](mailto:support@loadbalancer.org).

7. Select the required **Authentication type** - either a **Password** or an **SSH Public key** can be used.

- If **SSH Public Key** is Selected:

- Specify a suitable **Username**.



- Select the *SSH public key source*.
- Specify the *SSH Key Type*.
- Specify the *Key pair name*.

▪ If *Password* is selected:

- Enter an appropriate *Username* and *Password*.

8. Click **Next : Disks >**.

### Configure Disks

#### VM disk encryption

Azure disk storage encryption automatically encrypts your data stored on Azure managed disks (OS and data disks) at rest by default when persisting it to the cloud.

Encryption at host



**i** Encryption at host is not registered for the selected subscription.  
[Learn more about enabling this feature](#)

#### OS disk

OS disk size

Image default (30 GiB)



OS disk type \*

Premium SSD (locally-redundant storage)



Delete with VM



Key management

Platform-managed key



Enable Ultra Disk compatibility



#### Data disks for LB1

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

LUN	Name	Size (GiB)	Disk type	Host cache...	Delete with VM
-----	------	------------	-----------	---------------	----------------

[Create and attach a new disk](#) [Attach an existing disk](#)

1. Select the required *OS disk size* and *OS disk type* - the defaults are appropriate for most deployments.



**Note** Information on the various disk types available in Azure can be found [here](#). Comparative disk pricing is available [here](#).

2. Click **Next : Networking >**.

### Configure Networking



## Network interface

When creating a virtual machine, a network interface will be created for you.

Virtual network *	<input type="text" value="QS-RG1-VNET1"/> <a href="#">Create new</a>
Subnet *	<input type="text" value="Public-Subnet (10.1.3.0/24)"/> <a href="#">Manage subnet configuration</a>
Public IP	<input type="text" value="(new) LB1-ip"/> <a href="#">Create new</a>
NIC network security group	<input type="radio"/> None <input type="radio"/> Basic <input checked="" type="radio"/> Advanced <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"><p><span style="color: #0078d4;">i</span> This VM image has preconfigured NSG rules</p><p><span style="color: #0078d4;">i</span> The selected subnet 'Public-Subnet (10.1.3.0/24)' is already associated to a network security group 'VNET1-Public-Subnet-nsg'. We recommend managing connectivity to this virtual machine via the existing network security group instead of creating a new one here.</p></div>
Configure network security group *	<input type="text" value="VNET1-Public-Subnet-nsg"/> <a href="#">Create new</a>
Delete public IP and NIC when VM is deleted	<input type="checkbox"/>
Enable accelerated networking	<input type="checkbox"/> The selected image does not support accelerated networking.

## Load balancing

You can place this virtual machine in the backend pool of an existing Azure load balancing solution. [Learn more](#)

Load balancing options	<input checked="" type="radio"/> None <input type="radio"/> Azure load balancer Supports all TCP/UDP network traffic, port-forwarding, and outbound flows. <input type="radio"/> Application gateway Web traffic load balancer for HTTP/HTTPS with URL-based routing, SSL termination, session persistence, and web application firewall.
------------------------	---

1. Configure the *Virtual network, Subnet & Public IP* settings according to your requirements.

2. Configure the *Network Security Group* settings according to your requirements.

### i Note

Microsoft recommends that where possible Network Security Groups are associated with subnets rather than individual interfaces since this simplifies management.

- If you choose to create a new NSG rather than selecting an existing one, the following inbound rules are included by default:



## Inbound rules ⓘ

1010: HTTP\_access

Any

Custom (TCP/9080)

1020: HTTPS\_access

Any

Custom (TCP/9443)

1030: default-allow-ssh

Any

SSH (TCP/22)

[+ Add an inbound rule](#)

- These inbound rules are required for managing the load balancer
- The rules can be edited by clicking the *Create new* link under the network security group dropdown
- Specify additional inbound rules for the ports used for your load balanced applications, e.g. TCP 80 and TCP 443 if you're load balancing web servers, TCP 3389 if you're load balancing RDP etc.
- To specify additional inbound rules, click **Add an inbound rule**. The example below shows additional ports TCP/80 and TCP/443 for load balanced web server traffic.

## Inbound rules ⓘ

1010: HTTP\_access

Any

Custom (TCP/9080)

1020: HTTPS\_access

Any

Custom (TCP/9443)

1030: default-allow-ssh

Any

SSH (TCP/22)

1040: Port80

Any

HTTP (TCP/80)

1050: Port443

Any

HTTPS (TCP/443)

[+ Add an inbound rule](#)

### Note

The rules can also be edited after the NSG is created.

3. Click **OK**.

4. Under the "Load Balancing" section, if you have already configured an Azure load balancer in preparation for configuring an HA pair, select the **Azure load balancer** option and select the required load balancer from the dropdown. Alternatively, the Azure load balancer can be created later, in which case leave the option set to **None**.



### Note

When deploying an HA pair of Loadbalancer.org instances, they must be deployed behind an Azure load balancer. For more information please refer to [High Availability in Azure](#).

5. Click **Next : Management >**.

### **Configure Management**

1. Configure the Management settings according to your requirements.
2. Click **Next : Monitoring >**.

### **Configure Monitoring**

1. Configure the Monitoring settings according to your requirements.
2. Click **Next : Advanced >**.

### **Configure Advanced**

1. Configure the Advanced settings according to your requirements.
2. Click **Next : Tags >**.

### **Configure Tags**

1. Configure the Tags settings according to your requirements.
2. Click **Next : Review & Create >**.

### **Review & Create**

1. Review all details, terms and settings, enter your **Name**, **Preferred e-mail address** and **Preferred phone number** in the fields provided and if you're happy to proceed, click **Create**.
2. If the **Authentication type** was set to **SSH Public Key** and a new key was added, you'll be prompted to save the key - click **Download private key and create resource**.

## **7. Accessing the Appliance**

### **7.1. Accessing the Appliance WebUI**

Using a browser, navigate to the public IP address or FQDN on port 9443:

**https://<Public IP address>:9443**

or

**https://<FQDN>:9443**

### Note

To configure an FQDN in Azure please refer to this [link](#).



 **Note**

You'll receive a warning about the WebUI's SSL certificate. This is due to the default self signed certificate that is used. If preferred, you can upload your own certificate - for more information, please refer to [Appliance Security Features](#).

 **Note**

If you need to change the port, IP address or protocol that the WebUI listens on, please refer to [Service Socket Addresses](#).

Log in to the WebUI using the following default credentials:

**Username:** loadbalancer

**Password:** loadbalancer

 **Note**

To change the password, use the WebUI option: [\*Maintenance > Passwords\*](#).

Once logged in, the WebUI is displayed:



Primary | Secondary      Active | Passive      Link      32 Seconds

**System Overview**

**Local Configuration**

**Cluster Configuration**

**Maintenance**

**View Configuration**

**Reports**

**Logs**

**Support**

**Live Chat**

**WARNING:** YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS.

Buy with confidence. All purchases come with a 90 day money back guarantee.  
Already bought? Enter your license key [here](#)

[Buy Now](#)

**System Overview** 2024-01-31 13:39:34 UTC

No Virtual Services configured.

**Network Bandwidth**

Bytes/s

3.0 k  
2.0 k  
1.0 k  
0.0

Tue 18:00      Wed 00:00      Wed 06:00      Wed 12:00

3k Min, 3k Avg, 608k Total  
0k Min, 0k Avg, 4k Total

**System Load Average**

System Load

1.0  
0.8  
0.6  
0.4  
0.2  
0.0

Tue 18:00      Wed 00:00      Wed 06:00      Wed 12:00

1m average: 0.37 Min, 0.37 Avg, 0.37 Max  
5m average: 0.13 Min, 0.13 Avg, 0.13 Max  
15m average: 0.05 Min, 0.05 Avg, 0.05 Max

**Memory Usage**

Bytes

4.0 G  
3.0 G  
2.0 G

Tue 18:00      Wed 00:00      Wed 06:00      Wed 12:00

Memory Usage

Bytes

4.0 G  
3.0 G  
2.0 G

Tue 18:00      Wed 00:00      Wed 06:00      Wed 12:00

### 7.1.1. WebUI Menu Options

**System Overview** - Displays a graphical summary of all VIPs, RIPv and key appliance statistics

**Local Configuration** - Configure local host settings such as IP address, DNS, system time etc.

**Cluster Configuration** - Configure load balanced services such as VIPs & RIPv

**Maintenance** - Perform maintenance tasks such as service restarts and taking backups

**View Configuration** - Display the saved appliance configuration settings

**Reports** - View various appliance reports & graphs

**Logs** - View various appliance logs

**Support** - Create a support download, contact the support team & access useful links

**Live Chat** - Start a Live Chat session with one of our Support Engineers

## 7.2. Appliance Security



For full details of each security mode and all other security related features, please refer to



### 7.2.1. Security Mode

To control how the appliance is accessed and which features are enabled, 3 security modes are provided:

- **Custom** - In this mode the security options can be configured to suit your requirements
- **Secure - (Default)** - In this mode:
  - All console access and SSH password access is disabled
  - WebUI connections are forced to use HTTPS
  - Access to the *Local Configuration > Execute shell command* menu option is disabled
  - The Firewall Script & the Firewall Lockdown Wizard Script cannot be edited
- **Secure - Permanent** - This mode is the same as **Secure** but once set it cannot be changed

**(①) Important** Setting the security mode to **Secure - Permanent** is irreversible.

To configure the Security Mode:

1. Using the WebUI, navigate to: *Local Configuration > Security*.
2. Select the required **Appliance Security Mode**. If **Custom** is selected, configure the required options.
3. Click **Update**.

### 7.2.2. Passwords

#### The loadbalancer WebUI account

The password for the **loadbalancer** WebUI user account is set to "loadbalancer" by default. This can be changed using the WebUI menu option: *Maintenance > Passwords*.

#### The root Linux account

It's not possible to directly log in as root. If root access is required, once you've logged into the console/SSH session using the credentials defined during instance deployment, run the following command:

```
$ sudo su
```

## 7.3. Appliance Software Update

For v8.6.0 and later, the appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. If an update is found, a message will be displayed at the top of the screen as shown in the following example:



**Information:** Update 8.13.2 is now available for this appliance.

**Online Update**

To start the update process click **Online Update**.

The update check can also be initiated manually.

To initiate an online update check:

1. Using the WebUI, navigate to: **Maintenance > Software Update**.
2. Click **Online Update**.

**Note**

Updates are incremental for versions prior to v8.8.0, so repeat the process ignoring calls to restart services or reboot the appliance until you reach v8.8.0. Once at v8.8.0, when the update is next triggered the appliance will be updated to the latest version in a single step.

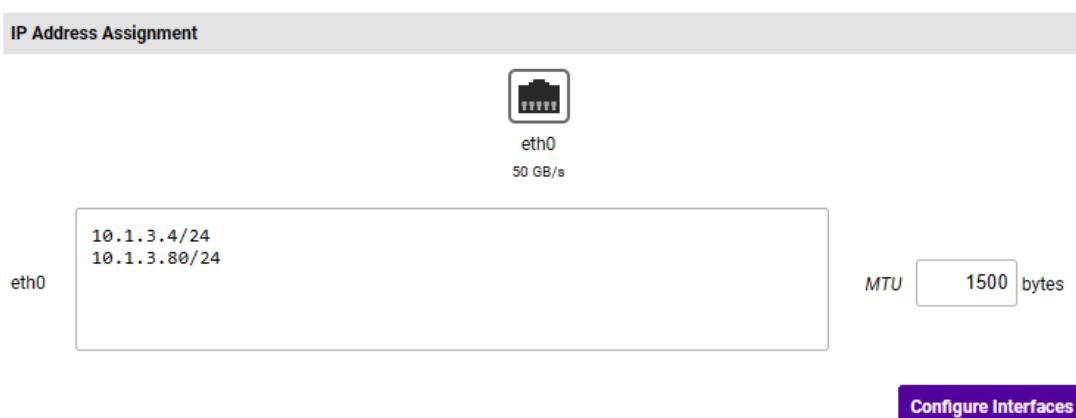
## 7.4. Appliance Licensing

If you've deployed the BYOL version of the appliance, by default it runs as a 30 day trial and is completely unrestricted during this time. After 30 days, the appliance continues to work but it's no longer possible to make changes to the configuration. When a license is purchased, you'll be provided with a license key file by our sales team. This must then be installed on your appliance. To install the license, use the WebUI option: **Local Configuration > License Key** to browse to and select the license file provided. Once selected, click **Install License Key** to apply the license. We recommend that you should check for updates **before** applying the license key.

## 7.5. Enterprise Azure Non-standard WebUI Menu Options

Enterprise Azure has some differences to the standard hardware/virtual product range due to the way the Microsoft Azure environment works. The menu options that work differently are detailed below. For all others please refer to our main [Administration Manual](#).

### 1) Local Configuration > Network Interface Configuration



IP Address Assignment

eth0

50 GB/s

10.1.3.4/24  
10.1.3.80/24

MTU 1500 bytes

Configure Interfaces



This menu option works in a very similar way to the standard product range, although please note the following:

- On initial deployment, a single IP private address is allocated (either static or dynamic depending on the chosen setting)
- Additional addresses can be added as shown above (10.1.3.80/24) - this is required when you require multiple VIPs on different IP addresses
- To add an additional IP address, enter the new address below the existing address as shown in the example above, then click **Configure Interfaces**

**(1) Important** If an IP address is added, you'll also need to add the same IP address to the Network Interface on the load balancer VM via the Azure portal. If this is not done, Azure will not be aware of the new address.

**(1) Important** If the IP address allocated to the VM on initial deployment (normally the first in the list) is changed, make sure that you also update the IP address via the Azure Portal.

## 7.6. Accessing the Appliance using SSH

When the appliance is deployed, **Authentication type** must be set to either **SSH Public key** or **Password**. When set to **SSH Public Key**, a new a key can be created or an existing key can be selected. The SSH key specified for the VM must be used with the SSH client machine to access the VM.

### 7.6.1. Accessing the Appliance from Linux

Start SSH specifying the private key file and login as the user defined when deploying the VM, e.g.

```
# ssh -i lb-privatekeyfile azureuser@1.2.3.4
```

### 7.6.2. Accessing the Appliance from Windows using PuTTY

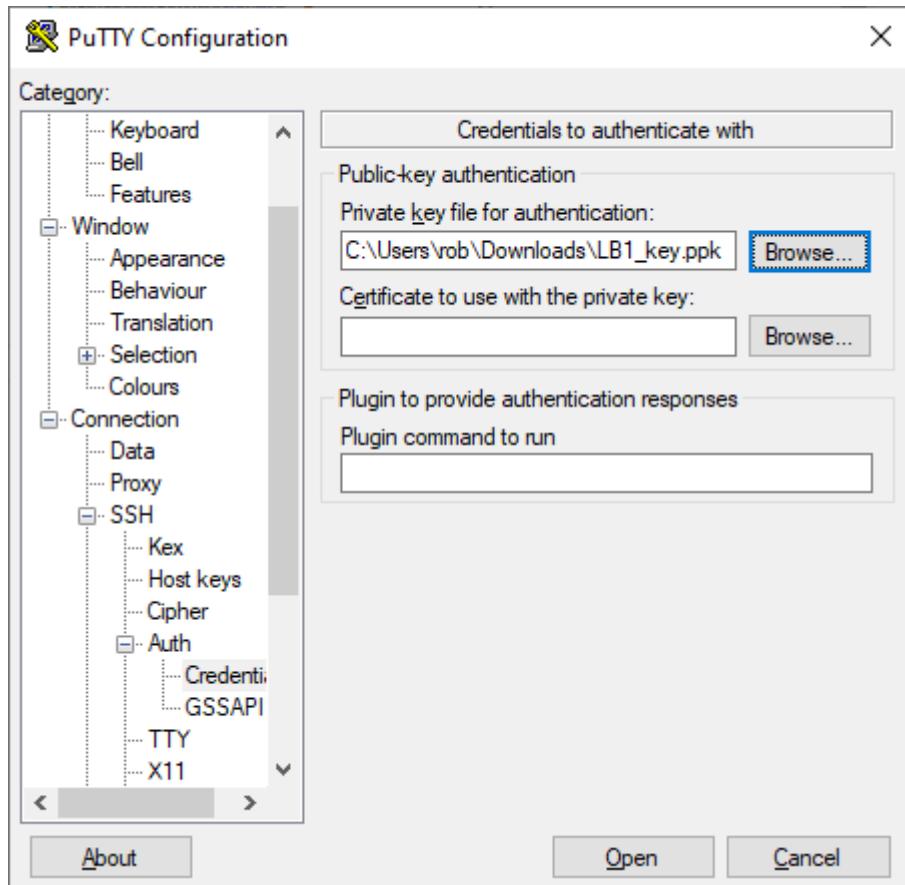
#### Step 1 - Convert the Private Key to Putty Format (ppk)

1. Run PuttyGen.
2. Click **Load** and browse to and select private key for the VM.
3. Click **OK** to the successful import message.
4. Specify a passphrase if required and click **Save Private Key**.
5. Specify a filename and **Save**.

#### Step 2 - Access the Appliance

1. Run PuTTY.
2. Expand the SSH section and select **Auth > Credentials** as shown below.





3. Click **Browse** and select the private key (.ppk) just created.
4. Click **Open** to start the SSH session.
5. Login using the username specified when deploying the instance, no password will be required.
6. Once logged in, full root access can be enabled using the following command:

```
$ sudo su
```

## 8. High Availability in Azure

We recommend that 2 appliances are deployed as an HA clustered pair to avoid introducing a single point of failure. In Azure, you should configure your HA pair **FIRST** before setting up your load balanced services.

### 8.1. Key Concepts

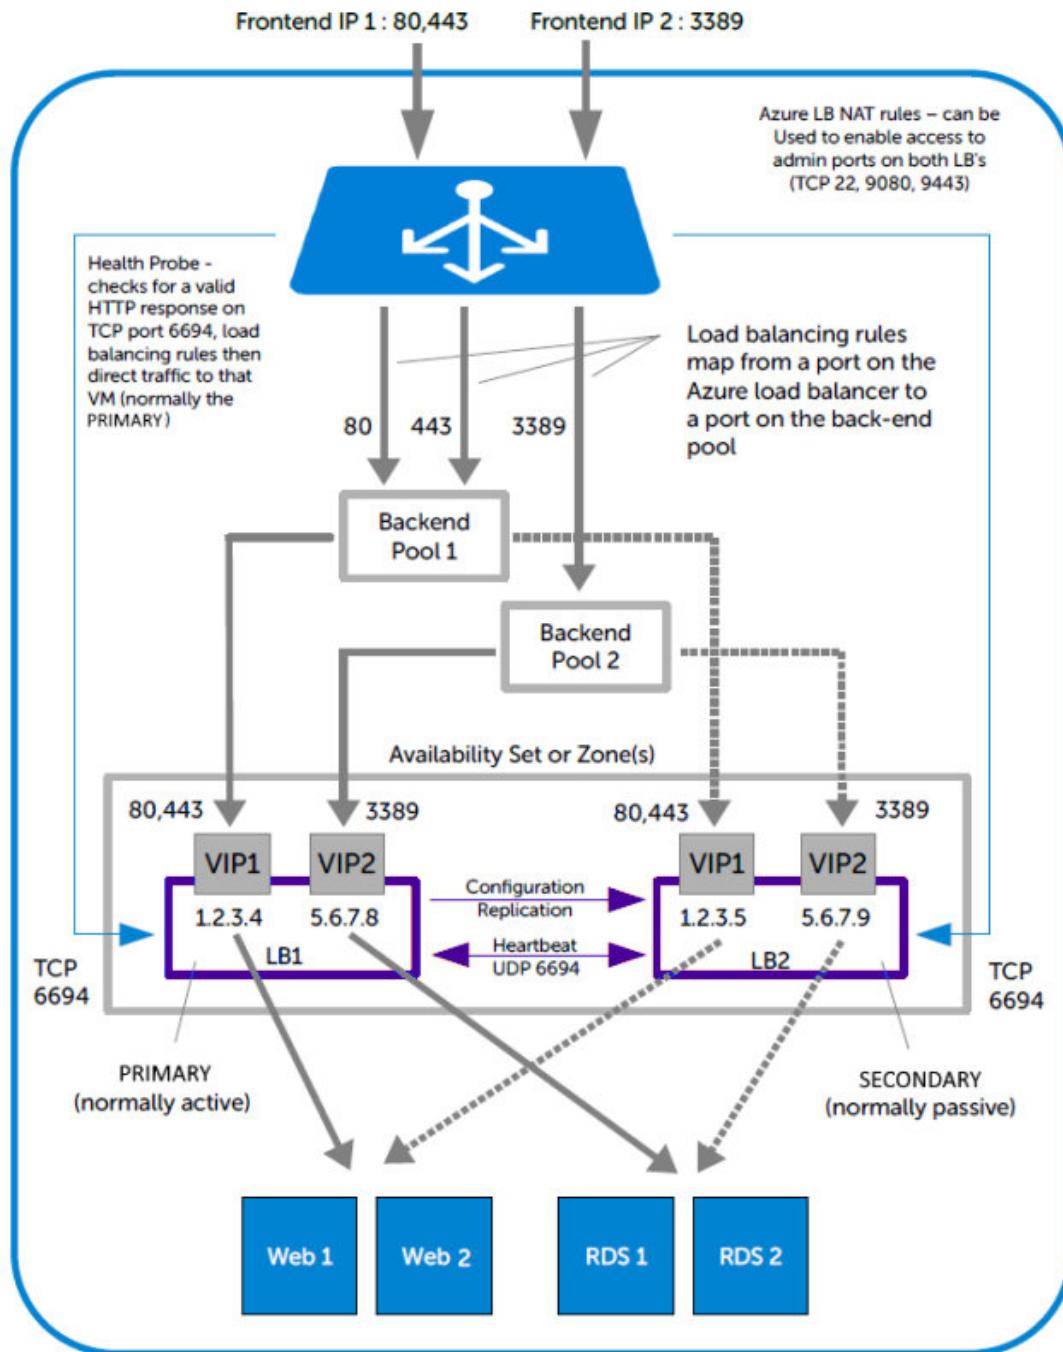
1. In our standard hardware / virtual product, when a failover occurs, the **same** VIP address is brought up on the passive device. In Azure, in order to minimize the time taken for the failover a different approach is used.
2. In Azure, when creating a Virtual Service on an HA pair, **2** private IPs must be specified - one to be used when the Primary is active and one to be used when the Secondary is active.
3. An Azure load balancer is used in front of the Loadbalancer.org HA pair to direct inbound traffic to the **active** appliance.
4. Both Primary & Secondary appliances must be in the same Availability Set or deployed within the same



Availability Zone or split across 2 different Availability Zones.

## 8.2. Implementing HA in Azure

The following diagram shows how HA is configured in Azure. As shown, two Loadbalancer.org VMs are configured as a clustered pair in combination with an Azure load balancer.



- LB1 and LB2 are configured as an HA pair. In this mode, one device is active (typically the Primary appliance) and the other is passive (typically the Secondary appliance).
- The private IPs for the VIP on the Primary & Secondary are selected using drop-downs within the VIP configuration screen. These drop-downs are only displayed once the pair is configured. They are populated with the IPs that are assigned to the network interface using the WebUI option: *Local Configuration > Network Interface Configuration*.



- The probe service on TCP port 6694 is up on the active appliance (LB1) and down on the passive appliance (normally LB2). The active appliance responds with **200 OK**.
- The Azure load balancer probes port 6694 on LB1 and LB2 and then forwards traffic to the active load balancer appliance (normally LB1).
- If the Primary appliance fails for any reason, the passive appliance will detect this, become active and bring up the probe service on port 6694. In turn, the Azure load balancer detects this and will then forward traffic to the Secondary device (LB2).
- If your configuration includes VIPs with multiple ports or if you have multiple VIPs you'll need to setup multiple *Load balancing rules* to map from the Azure load balancer's Frontend IP to the appropriate *Backend Pool* and appropriate port. Also, you may need to setup multiple *Frontend IP Configurations & Backend Pools* depending on whether your VIPs share the same IP or have unique IP addresses, and whether the load balanced servers are common between VIPs or unique. The same *Health-probe* should be used for all *Load balancing rules*.

## 9. Configuration Examples

This section presents 2 example configurations that illustrate how the load balancer is deployed. Web servers are used in the examples, although the same concepts apply to other applications.

### 9.1. Example 1 - Load Balancing Web Servers, HA Configuration, 1 Subnet, Layer 7

This example demonstrates how to configure an HA pair of load balancers and then configure a layer 7 VIP to load balance 2 web servers. The Loadbalancer.org instances are deployed across 2 Availability Zones and a Standard SKU Azure load balancer is used to route traffic to the active appliance.

#### Step 1 - Deploy VMs

1. Deploy 2 load balancer instances - one to be the Primary, the other the Secondary as described in [Deploying Enterprise Azure From the Marketplace](#). Ensure that each load balancer is deployed in a different Availability Zone.
2. Deploy the web server VMs into the same subnet.

#### Step 2 - Configure Network Security Group Settings

1. Ensure that your Network Security Group(s) permit the following communication between the 2 VMs:
  - TCP port 22 (SSH)
  - UDP port 6694 (heartbeat)
  - ICMP Ping (Heartbeat)

 **Note**

These requirements are covered by default within the same Virtual Network. Please refer to [this link](#) for more information on default rules.



2. Ensure that your Network Security Group(s) permit the following inbound communication from the Azure load balancer to both VMs:
  - TCP port 6694 (Azure load balancer health probe)

 **Note**

This requirement is covered by default within the same Virtual Network. Please refer to [this link](#) for more information on default rules.

3. Ensure that your Network Security Group(s) permit inbound traffic on port 80.

### Step 3 - Add the IP Address to be used for the VIP to the Primary & Secondary VMs using the Azure Portal

1. In the Azure Portal select *Virtual Machines*.
2. Select the Primary VM.
3. Select *Networking > Network Settings*, then click the Network Interface.
4. Select *IP Configurations*.
5. Click **Add**.

Name \*

IP version

Type

#### Private IP address settings

Allocation  Dynamic

Static

Private IP address \*

#### Public IP address settings

Associate public IP address

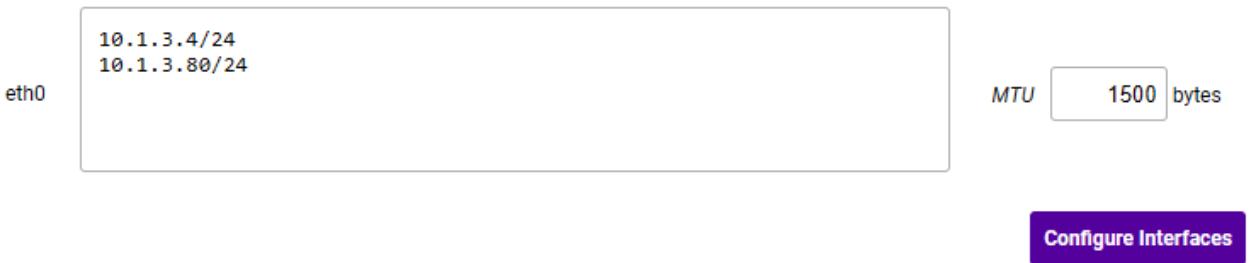
- Enter a suitable name for the IP address, e.g. **WebCluster**
- Set *Private IP address Allocation* to **Static**
- Enter the IP address to be used for the VIP, e.g. **10.1.3.80**
- Click **Add**

6. Now repeat steps 1-5 on the Secondary VM for the VIP using a corresponding (but different) IP address, e.g. **10.1.3.81**.



## Step 4 - Add the IP Address to be used for the VIP to the Primary & Secondary VMs using the Appliance WebUI

1. On the Primary, navigate to the WebUI option: *Local Configuration > Network Interface Configuration*



- Add the IP you intend to use for the VIP on the Primary appliance - use CIDR notation, e.g. **10.1.3.80/24**

2. Now repeat this step to add the IP you intend to use for the VIP on the Secondary appliance - use CIDR notation, e.g. **10.1.3.81/24**

## Step 5 - Add & Configure the Azure Load Balancer

1. In the Azure Portal select *Load balancers*.

2. Click **Create**

The screenshot shows the 'Create Load Balancer' wizard in the Azure Portal. In the 'Subscription' section, 'Loadbalancer.org Pay-As-You-Go' is selected. In the 'Resource group' section, 'QS-RG1' is selected. The 'Instance details' section includes fields for 'Name' (set to 'AzureLB'), 'Region' (set to 'UK South'), 'SKU' (set to 'Standard'), 'Type' (set to 'Public'), and 'Tier' (set to 'Regional'). Each of these fields has a green checkmark indicating they are valid.

- Configure the *Subscription & Resource group* settings according to your requirements
- Enter a suitable *Name* for the instance, e.g. **AzureLB**
- Select the required *Region*
- Leave the Set *SKU* to **Standard**
- If deploying within a private network set *Type* to **Internal**, if it's public facing select **Public**



- Leave the *Tier* set to **Regional**

3. Click **Next: Frontend IP configurations >**

4. Click **Add a Frontend IP configuration**.

Name *	<input type="text" value="WebClusterFrontendIPConfig"/>
IP version	<input checked="" type="radio"/> IPv4 <input type="radio"/> IPv6
IP type	<input checked="" type="radio"/> IP address <input type="radio"/> IP prefix
Public IP address *	<input type="text" value="(new) WebClusterPublicIP"/> <span style="font-size: small;">▼</span> <a href="#">Create new</a>
Gateway Load balancer	<input type="text" value="None"/> <span style="font-size: small;">▼</span>

- Configure the *Public IP address settings* (for external deployments) or the *Virtual Network* settings (for internal deployments) according to your requirements

5. Click **Save**.

6. Click **Next : Backend pools >**

7. Click **Add a backend pool**.

Name *	<input type="text" value="WebClusterBackendPool"/>
Virtual network	<input type="text" value="QS-RG1-VNET1 (QS-RG1)"/> <span style="font-size: small;">▼</span>
Backend Pool Configuration	<input checked="" type="radio"/> NIC <input type="radio"/> IP address

### IP configurations

IP configurations associated to virtual machines and virtual machine scale sets must be in same location as the load balancer and be in the same virtual network.

+ Add | X Remove

- Enter an appropriate *Name*.
- Select the *Virtual Network*.
- Set the *Backend Pool Configuration* to **NIC**.

8. Under *IP Configurations*, click **Add**.



Resource Name	Resource group	Type	IP configuration	IP Address	Availability set	Tags
<b>Virtual machine (4)</b>						
<input checked="" type="checkbox"/> LB1	QS-RG1	Virtual machine	WebCluster1	10.1.3.80	-	-
<input type="checkbox"/> LB1	QS-RG1	Virtual machine	ipconfig1	10.1.3.4	-	-
<input checked="" type="checkbox"/> LB2	QS-RG1	Virtual machine	WebCluster1	10.1.3.81	-	-
<input type="checkbox"/> LB2	QS-RG1	Virtual machine	ipconfig1	10.1.3.6	-	-

- Select the IP addresses that correspond to the VIP on each appliance (**10.1.3.80 & 10.1.3.81**) as shown above.

9. Click **Add**.

#### IP configurations

IP configurations associated to virtual machines and virtual machine scale sets must be in same location as the load balancer and be in the same virtual network.

<input type="button" value="Add"/>	<input type="button" value="Remove"/>					
Resource Name	Resource group	Type	IP configuration	IP Address	Availability set	
LB1	QS-RG1	Virtual machine	WebCluster1	10.1.3.80	-	
LB2	QS-RG1	Virtual machine	WebCluster1	10.1.3.81	-	

10. Click **Save**.

 **Note**

If you have multiple VIPs on different IPs you'll need to setup a Backend Pool for each of these. This is illustrated in the diagram in [Implementing HA in Azure](#).

11. Click **Next : Inbound rules >**.

12. Click **Add a load balancing rule**.



Name *	LbRuleWebCluster
IP version *	<input checked="" type="radio"/> IPv4 <input type="radio"/> IPv6
Frontend IP address * ⓘ	WebClusterFrontendPConfig (To be cre... ▾)
Backend pool * ⓘ	WebClusterBackendPool ▾
Protocol	<input checked="" type="radio"/> TCP <input type="radio"/> UDP
Port *	80 ▾
Backend port *	80 ▾
Health probe *	(new) LBprobe (TCP:6694) ▾
<a href="#">Create new</a>	
Session persistence	None ▾
<small> ⓘ Session persistence specifies that traffic from a client should be handled by the same virtual machine in the backend pool for the duration of a session. <a href="#">Learn more.</a></small>	

- Enter a suitable **Name**.
- Select the **IP Version**.
- Select the **Frontend IP address** created previously
- Select the **Backend pool** created previously
- Set the **Protocol** to **TCP**
- Set the **Port** to **80**
- Set the **Backend port** to **80**
- Under **Health Probe** click **Create new**



Name *	LBprobe
Protocol *	HTTP
Port * ⓘ	6694
Path * ⓘ	/
Interval (seconds) * ⓘ	5

- Enter a suitable **Name**
- Set the **Protocol** to **HTTP**
- Leave the **Path** set to its default value
- Set the **Port** to **6694**
- Leave the **Interval** at its default value
- Click **Save**
- leave **Session Persistence** set to **none** - session persistence is not required since the Azure Load balancer will simply send all traffic to the working Loadbalancer.org appliance, i.e the appliance that is responding with a **200 OK** to the HTTP probe on TCP port 6694

13. Click **Save**.

**ⓘ Note**

If your configuration includes other ports (e.g. HTTPS port 443) or if you have multiple VIPs you'll need to setup multiple **Load balancing rules** to map from the Azure load balancer's Frontend IP to the appropriate **Backend Pool** and appropriate port. Also, you may need to setup multiple **Frontend IP Configurations & Backend Pools** depending on whether your VIPs share the same IP or have unique IP addresses, and whether the load balanced servers are common between VIPs or unique. The same **Health-probe** should be used for all **Load balancing rules**. This is illustrated in the diagram in [High Availability in Azure](#).

14. Click **Next : Outbound rules >**.

15. Click **Next : Tags >**.

- Configure Tags according to your requirements

16. Click **Next : Review + Create >**

- Once validated, review the settings and click **Create**

## Step 6 - Configure the HA Clustered Pair

1. Open the WebUI on the Primary appliance.
2. Navigate to: *Cluster Configuration > High Availability Configuration*.



## Create a Clustered Pair



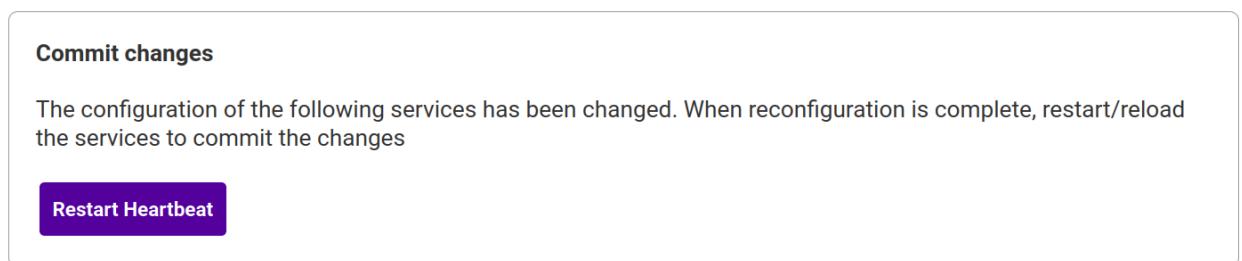
Local IP address  
10.1.3.4

IP address of new peer  
10.1.3.6

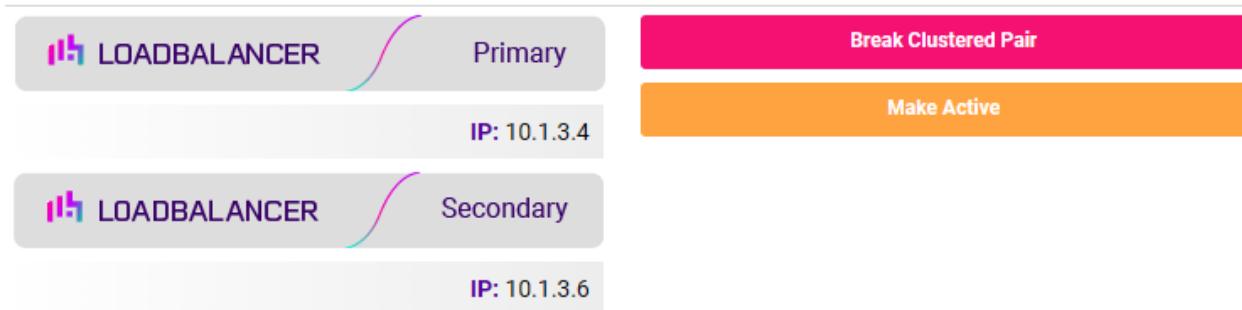
Password for *loadbalancer* user on peer  
.....

Add new node

3. In the *IP address of new peer* field, enter the Secondary appliance's private IP address.
4. In the *Password for loadbalancer user on peer* field enter the relevant password, the default password is "loadbalancer".
5. Click **Add new node**.
6. Once the pairing configuration has finished, any service restart messages and the confirmed pair message will be displayed as shown below:



## High Availability Configuration - primary



7. Restart the services using the buttons presented, in this case Heartbeat.

## Step 7 - Enable service control (IMPORTANT)

1. On the Primary appliance, navigate to: *Cluster configuration > Floating IPs*.



New Floating IP

10.1.3.100

Add Floating IP

2. In the **New Floating IP** field enter an unused IP address in the same subnet as the appliances - this address is not used for any traffic, it's required to allow service control on both Primary & Secondary appliances.



Note

The chosen IP address should not be in use anywhere else in the deployment.

3. Click **Add Floating IP**.

## Step 8 - Configure the Virtual Service (VIP)

1. On the Primary appliance, navigate to: *Cluster Configuration > Layer 7 - Virtual Services* and click **Add a New Virtual Service**.
2. Enter the following details:

Virtual Service		[Advanced -]
Manual Configuration	<input type="checkbox"/>	
Create Backend Only	<input type="checkbox"/>	
Label	WebCluster	
IP Address	10.1.3.80	
Secondary Appliance IP Address	10.1.3.81	
Ports	80	
Protocol		[Advanced +]
Layer 7 Protocol	HTTP Mode	
Termination		
Create HAProxy SSL Termination	<input type="checkbox"/>	
		<b>Cancel</b> <b>Update</b>

3. Enter an appropriate label for the VIP, e.g. **WebCluster**.
4. Set the **IP Address** field to the IP address of the VIP when active on the *Primary* appliance (the same address as added earlier in steps 3 & 4), e.g **10.1.3.80**.
5. Set the **Secondary IP Address** field to the IP address of the VIP when active on the *Secondary* appliance (same address added earlier in steps 3 & 4), e.g **10.1.3.81**.



Note

To assign additional IP addresses to the appliance, use the WebUI option: *Local Configuration > Network Interface Configuration*. If an IP address is added, you'll also need



to add the same IP address to the Network Interface on the load balancer VM via the Azure portal.

6. Set the *Virtual Service Ports* field to **80**.
7. Leave *Layer 7 Protocol* set to **HTTP**.
8. Click **Update**.

### Step 9 - Configure the Real Servers (RIPs)

1. On the Primary appliance, navigate to: *Cluster Configuration > Layer 7 - Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	Web1	?
Real Server IP Address	10.1.3.11	?
Real Server Port	80	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	100	?
		<b>Cancel</b> <b>Update</b>

3. Enter an appropriate label for the RIP, e.g. **Web1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **10.1.3.11**.
5. Set the *Real Server Port* field to **80**.
6. Click **Update**.
7. Repeat the above steps to add additional Web Server(s).

### Step 10 - Finalizing the Configuration

To apply the new settings, HAProxy must be reloaded. This can be done using the buttons in the "Commit changes" box at the top of the screen or by using the *Restart Services* menu option:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.

### Step 11 - Verify synchronization state

1. Verify that the status on the Primary & Secondary is as follows:



*Primary Unit:*

**Primary | Secondary**

**Active | Passive**

**Link**

*Secondary Unit:*

**Primary | Secondary**

**Active | Passive**

**Link**



**Note**

For other possible states, please refer to Clustered Pair Diagnostics.

The Secondary can be made active by clicking **[Advanced]** in the information box that states "This device is currently passive", and then clicking the **Take over** button:

**System Overview**

2023-02-01 13:11:18 UTC

**Information:** This device is currently passive. Please see the active device for Virtual Service statistics.  
[ Advanced ]

## Step 12 - Testing

1. Browse to the Frontend IP address on the Azure load balancer on port 80 and verify that the web page is displayed.

## 9.2. Example 2 - Load Balancing Web Servers, Single Appliance, 2 Subnets, Layer 4 NAT Mode

This example uses 2 subnets - one public subnet for the load balancer and one private subnet for the web servers. The load balancer has a single network interface located in the first subnet. Routing rules for the second private subnet must be changed so that return traffic passes back via the load balancer. This is achieved by creating a custom routing table with the required rules, then associating this with the private subnet.



**Note**

This configuration is currently not supported in HA mode. In this mode, the custom routing rules would need to be dynamically modified to route via the Secondary appliance rather than the Primary if a failover occurs. This is currently not supported.

## Step 1 - Deploy VMs

1. Deploy the load balancer instance into the first (public) subnet as described in [Deploying Enterprise Azure From the Marketplace](#).
2. Deploy your required web server VMs into the second (private) subnet.

## Step 2 - Configure Network Security Group Settings



1. Ensure that your Network Security Group(s) permit the following inbound communication from the Azure load balancer to both VMs:

- TCP port 6694 (Azure load balancer health probe)

 **Note**

This requirement is covered by default within the same Virtual Network. Please refer to [this link](#) for more information on default rules.

2. Ensure that your Network Security Group(s) permit the following web server traffic:

- **Public Subnet:**

- Inbound – from 0.0.0.0/0 to port 80
  - Outbound – from 0.0.0.0/0 to private subnet on port 80

- **Private Subnet:**

- Inbound – from 0.0.0.0/0 to port 80

 **Note**

Where possible, always specify the minimum source IP range rather than 0.0.0.0/0 (all IPs). Bear in mind that NAT mode is source IP transparent so the source address of inbound packets will be the client.

### Step 3 - Add the IP Address to be used for the VIP using the Azure Portal

1. In the Azure Portal select *Virtual Machines*.
2. Select the load balancer VM.
3. Select *Networking > Network Settings*, then click the Network Interface.
4. Select *IP Configurations*.
5. Click **Add**.



Name *	WebCluster
IP version	IPv4
Type	Secondary

### Private IP address settings

Allocation	<input type="radio"/> Dynamic
	<input checked="" type="radio"/> Static
Private IP address *	10.1.3.80

### Public IP address settings

Associate public IP address	<input type="checkbox"/>
-----------------------------	--------------------------

- Enter a suitable name for the IP address, e.g. **WebCluster**
- Set *Private IP address Allocation* to **Static**
- Enter an appropriate IP address, e.g. **10.1.3.80**
- Click **Add**

### Step 4 - Configure a Custom Routing Table

1. Using the search option at the top of the page, search for "Route tables".
2. Click **Create**.

Subscription *	Loadbalancer.org Pay-As-You-Go
Resource group *	QS-RG1
	<a href="#">Create new</a>
<b>Instance details</b>	
Region *	UK South
Name *	RT1
Propagate gateway routes *	<input checked="" type="radio"/> Yes <input type="radio"/> No

3. Configure the *Subscription*, *Resource group* & *Region* settings according to your requirements.



4. Enter a suitable name for the Route table, e.g. **RT1**.
5. Click **Next**.
6. Define any required tags.
7. Click **Review + Create**.
8. Click **Create**.
9. Once created, select the newly created Route table.
10. Click **Routes** under **Settings** then click **Add**.

Route name \*

Destination type \* ⓘ

Destination IP addresses/CIDR ranges \* ⓘ

Next hop type \* ⓘ

Next hop address \* ⓘ

**ⓘ** Ensure you have IP forwarding enabled on your virtual appliance. You can enable this by navigating to the respective network interface's IP address settings.

11. Enter a suitable name for the route, e.g. **R1**.
12. Set the **Destination Type** to **IP Addresses**.
13. Set the **Destination \_IP addresses/CIDR ranges** to **0.0.0.0/0**.
14. Set the **Next hop type** to **Virtual appliance**.
15. Set the next hop address to the IP address of the load balancer in the public subnet, e.g. **10.1.3.4**.
16. Click **Add**.

**ⓘ Note**

As mentioned in the note in the screen shot above, IP forwarding must be enabled for the load balancer VM. This is covered below in step 5.

17. Click **Subnets** under **Settings**.
18. Click **Associate**.
19. Select the relevant **Virtual network** and set the **Subnet** to the private subnet.

Virtual network ⓘ

QS-RG1-VNET1 (QS-RG1) ▼

Subnet \* ⓘ

Private-Subnet ▼

20. Click **OK**.

### Step 5 - Enable IP Forwarding for the Load balancer VM

1. In the Azure Portal main menu, select *Virtual Machines*.
2. Select the Load balancer VM and click *Networking > Network Settings*.
3. Click the Network Interface for the VM.
4. Click *IP Configurations*.
5. Ensure that *IP forwarding* is enabled as shown below:

#### IP Settings

Enable IP forwarding ⓘ

6. Click **Apply**.

### Step 6 - Configure the Virtual Service (VIP)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services* and click **Add a New Virtual Service**.
2. Enter the following details:

Virtual Service		
Label	WebCluster	?
IP Address	10.1.3.80	?
Ports	80	?
Protocol		
Protocol	TCP	?
Forwarding		
Forwarding Method	NAT	?
		<span style="border: 1px solid #ccc; padding: 2px 10px; background-color: #f2f2f2; border-radius: 5px; margin-right: 5px;">Cancel</span> <span style="background-color: #592968; color: white; border: 1px solid #592968; padding: 2px 10px; border-radius: 5px; font-weight: bold;">Update</span>



3. Enter an appropriate label for the VIP, e.g. **WebCluster**.
4. Set the **Virtual Service IP Address** field to an appropriate value, e.g. **10.1.3.80**.
5. Set the **Virtual Service Ports** field to **80**.
6. Leave **Protocol** set to **TCP**.
7. Set the **Forwarding Method** to **NAT**.
8. Click **Update**.

### Step 7 - Configure the Real Servers (RIPS)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	Web1	?
Real Server IP Address	10.1.3.11	?
Real Server Port	80	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate label for the RIP, e.g. **Web1**.
4. Change the **Real Server IP Address** field to the required IP address, e.g. **10.1.3.11**.
5. Set the **Real Server Port** field to **80**.
6. Click **Update**.
7. Repeat the above steps to add additional Web Server(s).

#### Note

If you want your Real Servers to be able to access the outside world, i.e. the Internet in a public facing deployment, outbound requests passing via the load balancer must be NAT'd so that the source IP becomes the load balancer's own external address. This can be configured using the WebUI menu option: *Cluster Configuration > Layer 4 Advanced Configuration* and setting the **Auto-NAT** drop-down to **eth0**. You'll also need to modify the NSG for the public subnet to allow inbound requests from the private subnet to the external destination addresses/ports.

### Step 8 - Assigning a Public IP Address

For public facing deployments, you'll need to associate a public IP address with the private IP address used for



the VIP.

1. Select the load balancer VM in the Azure Portal.
2. Click **Networking > Network Settings**.
3. Select the Network Interface.
4. Select **IP Configurations**.
5. Click the IP configuration for the VIP.
6. Change Public IP address to **Enabled**.
7. Select an existing available Public IP address or create a new one.
8. Click **Save**.

#### Step 9 - Testing

1. Browse to the IP address associated with the VIP on port 80 and verify that the web page is displayed.

## 10. Testing & Verification

For additional guidance on diagnosing and resolving any issues you may have, please also refer to [Diagnostics & Troubleshooting](#).

## 11. More Information

Please refer to our website for all the latest [Manuals](#) and [Deployment Guides](#).

## 12. Loadbalancer.org Technical Support

Our highly experienced Support Engineers are on hand to help 24 hours a day, 365 days a year.

### 12.1. Contacting Support

If you have any questions regarding the appliance or need assistance with load balancing your application, please don't hesitate to contact [support@loadbalancer.org](mailto:support@loadbalancer.org).





**Visit us:** [www.loadbalancer.org](http://www.loadbalancer.org)

**Phone us:** +44 (0)330 380 1064

**Phone us:** +1 833 274 2566

**Email us:** [info@loadbalancer.org](mailto:info@loadbalancer.org)

**Follow us:** [@loadbalancer.org](https://www.loadbalancer.org)

## About Loadbalancer.org

Loadbalancer.org's mission is to ensure that its clients' businesses are never interrupted. The load balancer experts ask the right questions to get to the heart of what matters, bringing a depth of understanding to each deployment. Experience enables Loadbalancer.org engineers to design less complex, unbreakable solutions - and to provide exceptional personalized support.

