



Enterprise Azure Quick Start Guide

v8.0.1

rev. 1.1.3

Copyright © 2002 – 2016 Loadbalancer.org, Inc



Table of Contents

Introduction.....	3
About Enterprise Azure.....	3
Main Differences to the Non-Cloud Product.....	3
Why use Enterprise Azure?.....	3
Azure Terminology.....	4
Azure Deployment Models.....	4
Getting Started.....	4
Accessing the Portal.....	4
Deployment Concepts.....	5
Azure Management.....	6
Deploying Enterprise Azure From the Marketplace.....	7
Accessing the Appliance.....	11
Accessing the Appliance using the WUI.....	11
WUI Menu Options.....	12
Accessing the Appliance using SSH.....	13
Configuring the Appliance for SSH Authentication.....	13
Creating SSH Keys.....	13
Using Linux.....	13
Using Windows.....	14
Accessing the Appliance from Windows using PuTTY.....	14
Accessing the Appliance using SCP.....	18
Using Linux.....	18
Using Windows.....	18
Configuration Examples.....	19
1) Load Balancing Web Servers - Single Subnet, Layer 7.....	19
a) Setting up Azure.....	19
b) Setting up the Virtual Service.....	19
c) Setting up the Real Servers.....	20
d) Applying the new Layer 7 Settings.....	20
e) Testing & Verification.....	20
2) Load Balancing Web Servers - Single Subnet Layer 7 with SSL Termination.....	21
a) Setting up Azure.....	21
b) Setting up the Virtual Service.....	21
c) Setting up the Real Servers.....	22
d) Configuring SSL Termination.....	22
e) Applying the new Settings.....	24
f) Testing & Verification.....	24
Testing & Validation.....	25
Testing Load Balanced Services.....	25
Diagnosing VIP Connection Problems.....	25
Taking Real Servers Offline.....	26
Using Reports & Log Files.....	27
Loadbalancer.org Technical Support.....	27
Appendix.....	28
1. Updating the Agent when Using newer Instance Sizes.....	28
2. Company Contact Information.....	29

Introduction

Azure is Microsoft's cloud platform. It's a growing collection of integrated services– compute, storage, data, networking, and application. 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.

About Enterprise Azure

The core software is based on customized versions of Centos 6 / RHEL 6, Linux 3.10, HA-Linux, HAProxy, Pound & STunnel.

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.

Main Differences to the Non-Cloud Product

- The appliance can only have a single IP address, all configured virtual services must also use this address
- The network setup is customized for Microsoft Azure deployment
- Layer 4 NAT mode and Layer 4 DR mode are not supported, only Layer 7 mode is available
- The WUI is not accessible on HTTP port 9080, only HTTPS port 9443

Why use Enterprise Azure?

Microsoft enables users to configure a load balancer to load balance multiple Azure instances running in the cloud. This does provide basic load balancing functionality but is limited in several areas.

Loadbalancer.org's Enterprise Azure load balancer provides the following additional features & advantages:

1. Supports comprehensive Layer 7 load balancing functionality
2. Load balances both Azure based and non-Azure based servers
3. Supports Round Robin and Least Connection connection distribution algorithms
4. Supports customizable timeouts for custom applications beyond those offered by Azure
5. Supports comprehensive back-end server health-check options
6. Enables fallback servers to be configured and invoked when all load balanced servers/services fail
7. Provides extensive real time and historical statistics reports
8. Supports session distribution based on actual server load (utilizing Loadbalancer.org's feedback agent which is available for both Linux & Windows)
9. Supports SSL Termination
10. Supports RDP Cookie based persistence
11. Supports full integration with Remote Desktop Services Connection Broker

Azure Terminology

<u>Acronym</u>	<u>Terminology</u>
Windows Azure	Microsoft's Cloud platform
ARM	Azure Resource Manager - the latest Azure deployment model
Classic	The original Azure deployment model
BLOB	Binary Large Object for storing large data items, basically a large file
VNet	Virtual Network for grouping services

Azure Deployment Models

The Azure platform is currently in transition from the older Classic or Service Management Model, to the new Resource Manager deployment model. Resource Manager is designed to eventually replace the classic deployment model, which has been the model up to now for deploying virtual machine-based workloads in Azure.

The Resource Manager deployment model provides a new way to deploy and manage the services that make up your application. This new model contains important differences from the classic deployment model, and the two models are not completely compatible with each other. To simplify the deployment and management of resources, Microsoft recommends that you use Resource Manager for new resources, and, if possible, re-deploy existing resources through Resource Manager.

Resource Manager Model : This is the newest deployment model for Azure resources. Most newer resources already support this deployment model and eventually all resources will.

Classic (aka Service Manager) Model : This is the original model and is supported by most existing Azure resources today. New resources added to Azure will not support this model.

Getting Started

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://account.windowsazure.com/>

Accessing the Portal

Accessing the old portal : <https://manage.windowsazure.com>

Accessing the new portal : <https://portal.azure.com>

Deployment Concepts

The table below shows the main differences between the deployment models:

Item	Azure Classic	Azure Resource Manager
Cloud Service for Virtual Machines	Cloud Service was a container for holding the virtual machines that required Availability from the platform and Load Balancing.	Cloud Service is no longer an object required for creating a Virtual Machine using the new model.
Availability Sets	Availability to the platform was indicated by configuring the same "AvailabilitySetName" on the Virtual Machines. The maximum count of fault domains was 2.	Virtual Machines that require high availability must be included in the Availability Set. The maximum count of fault domains is now 3.
Load Balancing	Creation of a Cloud Service provides an implicit native Azure load balancer for the Virtual Machines deployed.	Native Azure Load balancers must be explicitly defined and can be internal or external.
Virtual IP Address (VIP)	Cloud Services will get a default VIP (Virtual IP Address) when a VM is added to a cloud service. The Virtual IP Address is the address associated with the implicit Azure load balancer.	Public IP Address can be Static (Reserved) or Dynamic. Dynamic Public IPs can be assigned to a Load Balancer. Public IPs can be secured using Security Groups.
Reserved IP Address	You can reserve an IP Address in Azure and associate it with a Cloud Service to ensure that the IP Address is sticky.	Public IP Address can be created in "Static" mode and it offers the same capability as a "Reserved IP Address". Static Public IPs can only be assigned to a Load balancer right now.
Public IP (PIP) per VM	Public IP Addresses can also associated to a VM directly.	Public IP Address can be Static (Reserved) or Dynamic. However, only dynamic Public IPs can be assigned to a Network Interface to get a Public IP per VM right now.
Endpoints	Input Endpoints needed to be configured on a Virtual Machine to open up connectivity for certain ports. One of the common modes of connecting to virtual machines done by setting up input endpoints.	The concept of Endpoints no longer exists. Instead create a network security group. A network security group is a set of firewall rules that control traffic to and from your virtual machine.
DNS Name	A cloud service would get an implicit globally unique DNS Name. For example: mycoffeeshop.cloudapp.net	DNS Names are optional parameters that can be specified on a Public IP Address resource. The FQDN will be in the following format: <label>.<region>.cloudapp.azure.com
Network Interfaces	Primary and Secondary Network Interface and its properties were defined as network configuration of a Virtual machine.	Network Interfaces are separately defined and associated with a VM. The lifecycle of the Network Interface is not tied to a Virtual Machine.

Azure Management

Many management tasks can be carried out using the Portal. However, there are still a number of tasks that must be done via Azure PowerShell or the CLI.

Azure PowerShell is a powerful scripting environment that can be used to control and automate the deployment and management of workloads in Azure.

The Azure CLI provides a set of open source, cross-platform commands for working with the Azure Platform.

Useful related Microsoft Links:

How to install and configure Powershell:

<https://azure.microsoft.com/en-gb/documentation/articles/powershell-install-configure/>

How to install and configure Azure CLI:

<https://azure.microsoft.com/en-gb/documentation/articles/xplat-cli/>

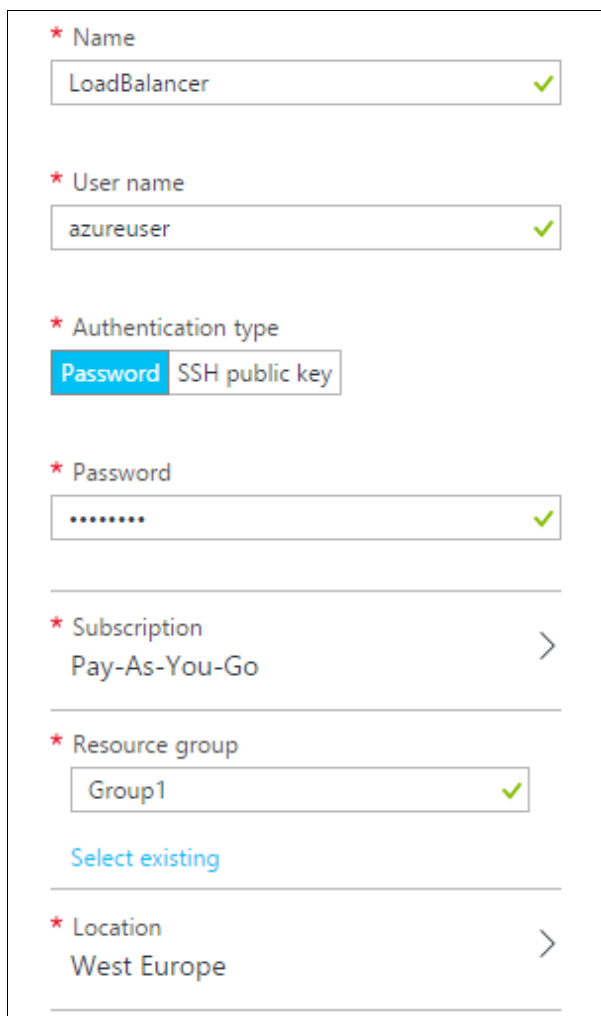


IMPORTANT NOTE : The load balancer has a single public IP address in Azure so all work-load and management services must be accessed via the same IP address.

Deploying Enterprise Azure From the Marketplace

(Using Resource Manager Model)

- Login to the Azure Management Portal: <https://portal.azure.com>
- In the Azure Management Portal select the **Virtual Machines** option
- Click **Add**
- In the search box type **Loadbalancer.org** and press <ENTER>
- Select *Loadbalancer.org Load Balancer for Azure*
- Set the deployment model to **Resource Manager**
- Click **Create**
- Define the required **Basic** settings:















The screenshot shows the 'Basic' configuration page for a Load Balancer in the Azure Marketplace. The form includes the following fields and options:

- Name:** A text input field containing 'LoadBalancer' with a green checkmark to its right.
- User name:** A text input field containing 'azureuser' with a green checkmark to its right.
- Authentication type:** A selection box with 'Password' selected (highlighted in blue) and 'SSH public key' as an alternative option.
- Password:** A text input field containing seven dots, with a green checkmark to its right.
- Subscription:** A dropdown menu showing 'Pay-As-You-Go' with a right-pointing chevron.
- Resource group:** A text input field containing 'Group1' with a green checkmark to its right. Below this field is a blue link labeled 'Select existing'.
- Location:** A dropdown menu showing 'West Europe' with a right-pointing chevron.

- Enter a suitable name, e.g. **LoadBalancer**
- Enter **azureuser** in the *User name* field

- Select the required *Authentication type*, a password or an SSH key can be used
N.B. Please refer to page 13 for more details on creating and using SSH keys
*If authentication is by Password, ensure the user name is set to **azureuser***
- Define a suitable password
- Set the *Subscription, Resource group* and *Location* according to your needs
- Click **OK**
- Select an appropriate size for your instance:

★ Recommended | [View all](#)

A1 Standard ★		A2 Standard ★		A3 Standard ★	
1	Core	2	Cores	4	Cores
1.75	GB	3.5	GB	7	GB
	2 Data disks		4 Data disks		8 Data disks
	2x500 Max IOPS		4x500 Max IOPS		8x500 Max IOPS
	Load balancing		Load balancing		Load balancing
	Auto scale		Auto scale		Auto scale
62.80		125.60		215.67	
GBP/MONTH (ESTIMATED)		GBP/MONTH (ESTIMATED)		GBP/MONTH (ESTIMATED)	

N.B. If you select a size other than A1, A2, or A3 Standard, you may need to manually update the Azure VM Agent. If a 'Deployment Failed' error is reported, please refer to page 28 in the Appendix for the steps to do this

- Click **Select**
- By default the load balancer will be allocated a Public IP address. If this is not required, under **Settings** click **Public IP address** and change the setting to **None**
- To configure access to the required ports, click **Network Security Group**
- A default Network Security Group with the same name as the load balancer will be displayed with the following settings:

* Name

LoadBalancer

Inbound rules ⓘ

1010: Web_interface	Any	Custom (TCP/9080)	✓	...
1020: Web_interface_ssl	Any	Custom (TCP/9443)	✓	...
1030: default-allow-ssh	Any	SSH (TCP/22)	✓	...

+ Add an inbound rule

Outbound rules ⓘ

No results.

+ Add an outbound rule

- Leave the default inbound rules, these are required for managing the load balancer
- Add additional inbound rules for the ports used for your load balanced services, e.g. TCP 80 and 443 if you are load balancing web servers , then click **OK**
- Configure any other **Settings** as required
- Click **OK**
- Check the Summary details are correct , e.g.

Basics	
Subscription	Pay-As-You-Go
Resource group	(new) Group1
Location	West Europe
Settings	
Computer name	LoadBalancer
User name	azureuser
Size	Standard A1
Disk type	Standard

Storage account	(new) group15968
Virtual network	(new) Group1
Subnet	(new) default (10.1.0.0/24)
Public IP address	(new) LoadBalancer
Network security group	(new) LoadBalancer
Availability set	None
Diagnostics	Enabled
Diagnostics storage account	(new) group15968

- Click **OK**
- Read the Purchase details and terms of use, and if happy to proceed click **Purchase**
- The load balancer will now be deployed, this will take a several minutes to complete
- Once deployed, the public IP address allocated will be displayed



NOTE : To enable full root access if required, the following command can be used once logged in to the appliance via ssh:

```
$ sudo su
```

For this to work without further modification, the username specified when setting up the appliance must be '**azureuser**' as mentioned above

Accessing the Appliance

Accessing the Appliance using the WUI

In a browser, navigate to the Public IP address or FQDN on port 9443 , i.e.

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

or

<https://<FQDN>:9443>

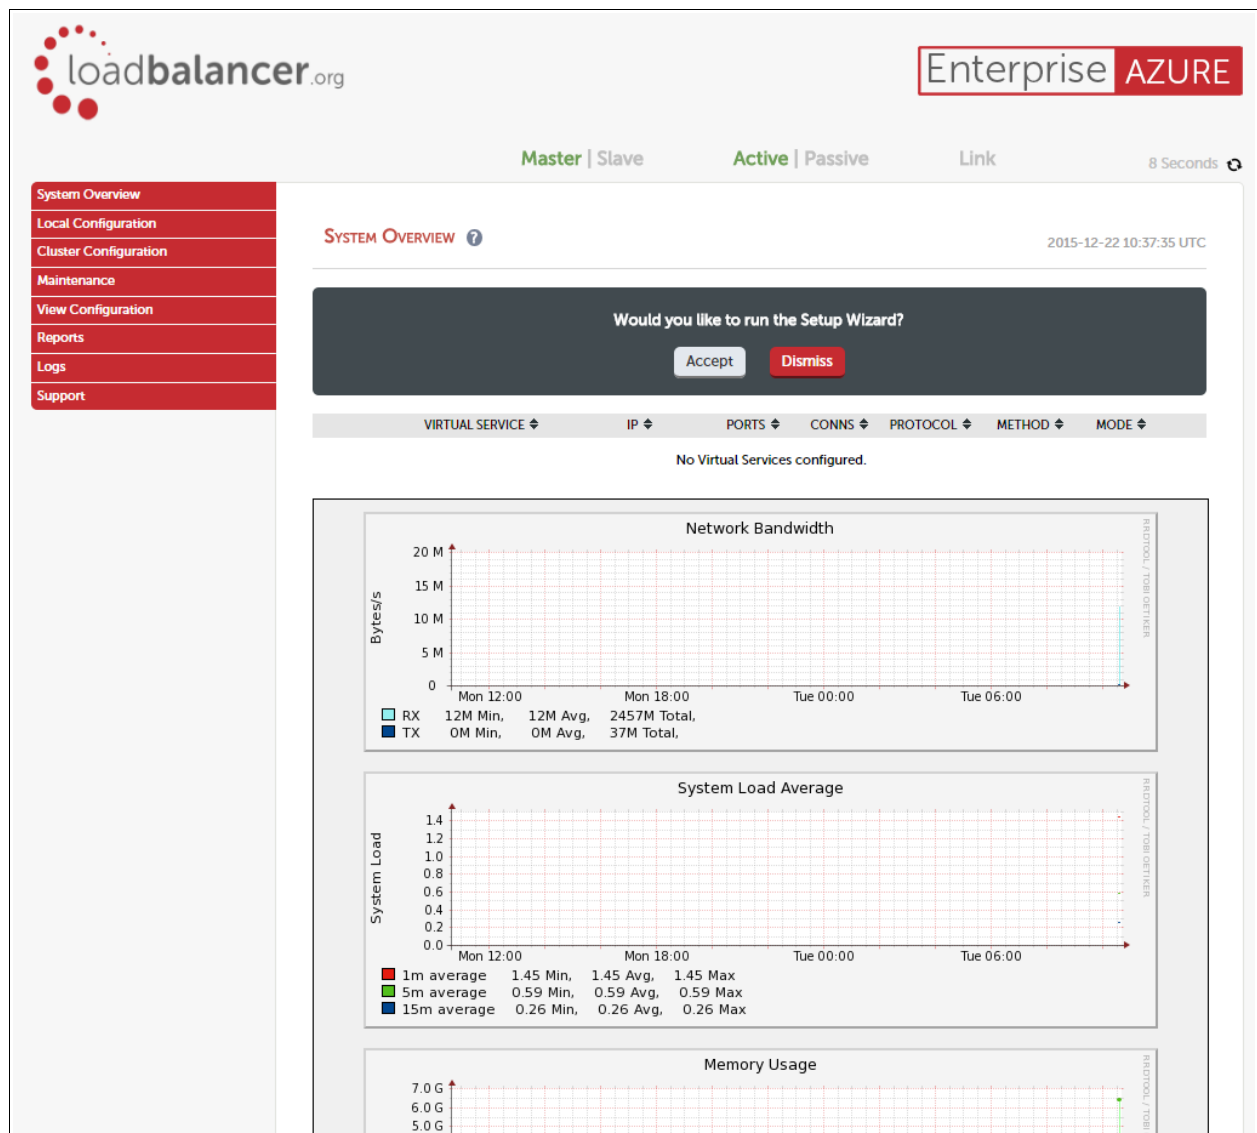
N.B. to configure an FQDN in Azure under the Resource Manager model please refer to [this link](#)

You'll receive a warning about the certificate as it's a self signed cert not related to an Internet based CA. Confirm you want to continue and a login prompt will be displayed. Use the following default credentials:

Username: *loadbalancer*

Password: *loadbalancer*

Once logged in, the following screen is displayed:



WUI Menu Options

The main menu options are as follows:

System Overview – Displays a graphical summary of all VIPs, RIPS and key appliance statistics

Local Configuration – Configure local host settings such as DNS, Date & Time etc.

Cluster Configuration – configure load balanced services such as VIPs & RIPs

Maintenance – Perform maintenance tasks such as service restarts and taking backups

View Configuration – Display the saved appliance configuration settings

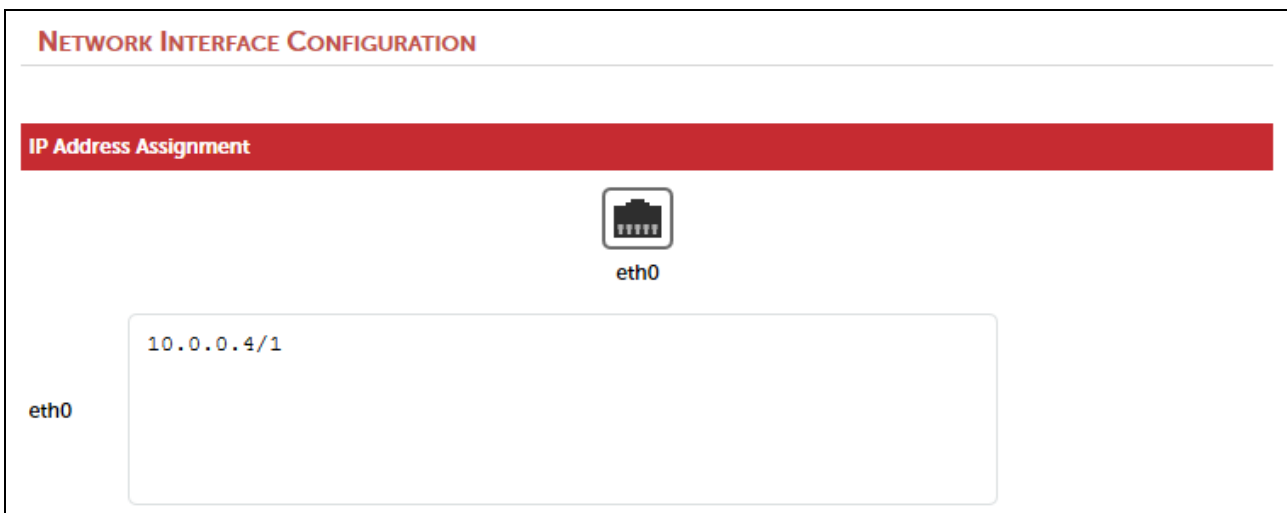
Reports – View various appliance reports & graphs

Logs – View various appliance logs

Support – Create a support download & contact the support team

The following sections detail the menu options that differ from our main product. For all others please refer to our main administration manual : <http://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf>

Local Configuration > Network Interface Configuration



Notes:

- Shows the private IP address allocated to the Virtual Machine, this cannot be changed via the WUI

Accessing the Appliance using SSH

Configuring the Appliance for SSH Authentication

Specify the PEM or .CER file when creating the Loadbalancer.org Enterprise Azure Instance. This is configured within the **Authentication type** section.

Creating SSH Keys

The current version of the Azure Management Portal only accepts SSH public keys that are encapsulated in an X509 certificate. The steps below show how to generate and use SSH keys with Azure under Linux and Windows.

Using Linux

STEP 1 - Install OpenSSL if not currently installed:

Under CentOS / Oracle Linux:

```
# yum install openssl
```

Under Ubuntu:

```
# apt-get install openssl
```

Under SLES & OpenSUSE:

```
# zypper install openssl
```

STEP 2 - Use OpenSSL to generate an X509 certificate with a 2048-bit RSA keypair:

(Answer the questions that OpenSSL prompts for (or you may leave them blank). The content in these fields is not used by the platform)

All Distros:

```
# openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout MyPrivateKey.key -out MyCert.pem
```

2 files are created:

- **MyPrivateKey.key** – this is used on the SSH client machine
- **MyCert.pem** – this is uploaded to the Loadbalancer.org Enterprise Azure Virtual Machine when the VM is deployed

STEP 3 - Change the permissions on the private key to secure it:

All Distros:

```
# chmod 600 MyPrivateKey.key
```

Using Windows

STEP 1 - Install OpenSSL if not currently installed:

Download and install OpenSSL light using one of the following links:

32 bit: https://slproweb.com/download/Win32OpenSSL_Light-1_0_2e.exe

64 bit: https://slproweb.com/download/Win64OpenSSL_Light-1_0_2e.exe

STEP 2 - Use OpenSSL to generate an X509 certificate with a 2048-bit RSA keypair:

(Answer the questions that the OpenSSL prompts for (or you may leave them blank). The content in these fields is not used by the platform)

```
C:/> openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout MyPrivateKey.key -out MyCert.pem
```

2 files are created:

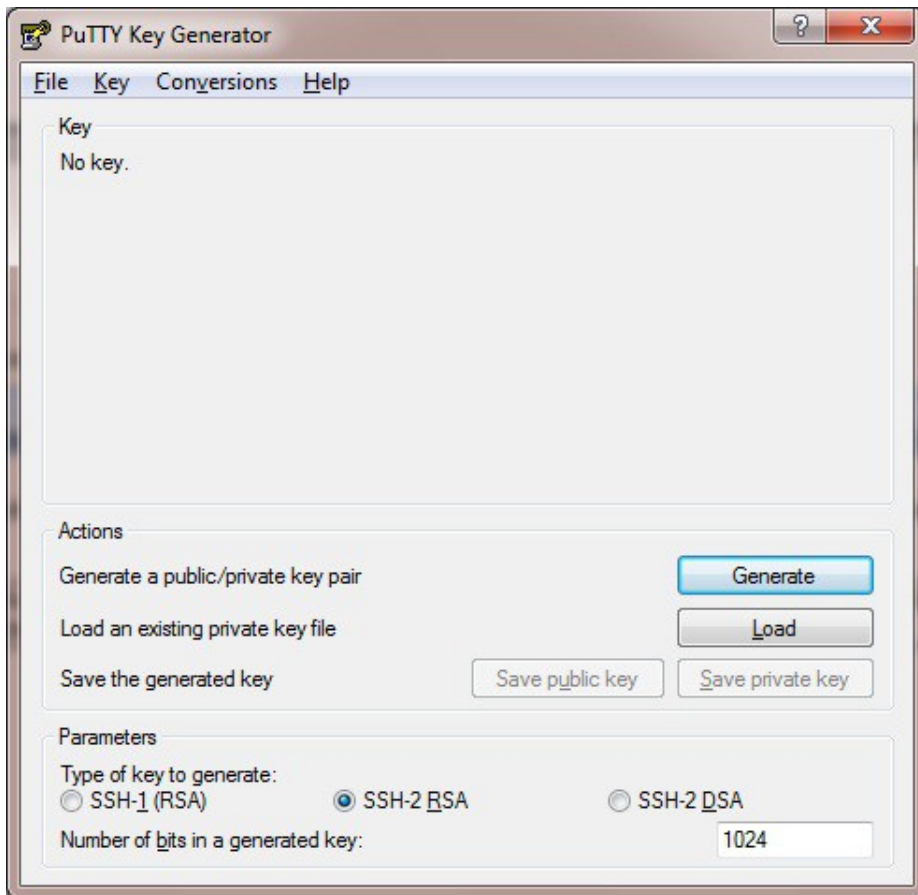
- **MyPrivateKey.key** – this is used on the SSH client machine
- **MyCert.pem** – this is uploaded to the Loadbalancer.org Enterprise Azure Virtual Machine when the VM is deployed

Accessing the Appliance from Windows using PuTTY

Download PuTTY from: <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

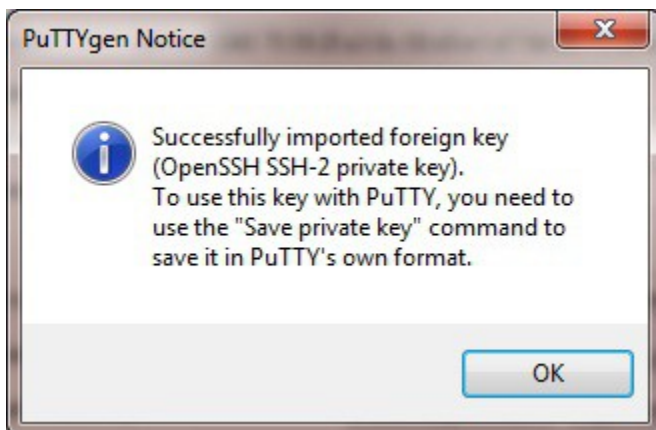
For PuTTY, the private key must be converted into an appropriate format. To do this the PuTTYgen utility (included with PuTTY) must be used.

Start PuTTYgen:

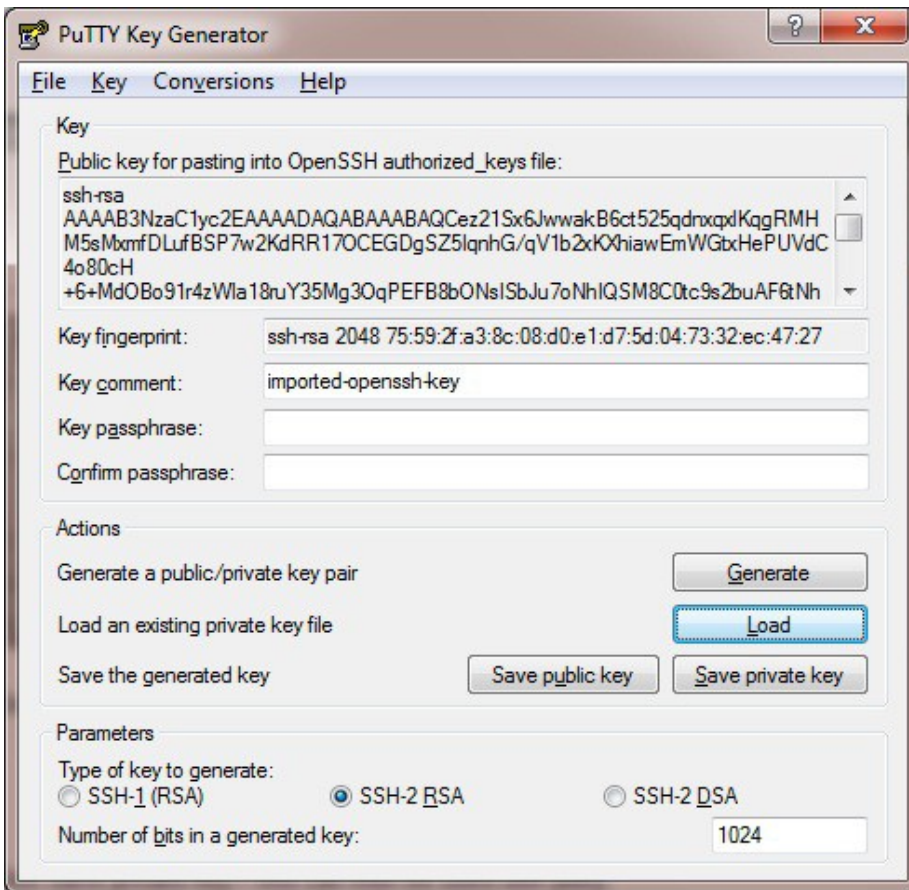


Click **Load**, change the file-type to all files and select the pem file saved earlier when creating your Key Pair.

You should see the following message:




Click **OK**



Now Click **Save private key** – this can then be used with PuTTY.

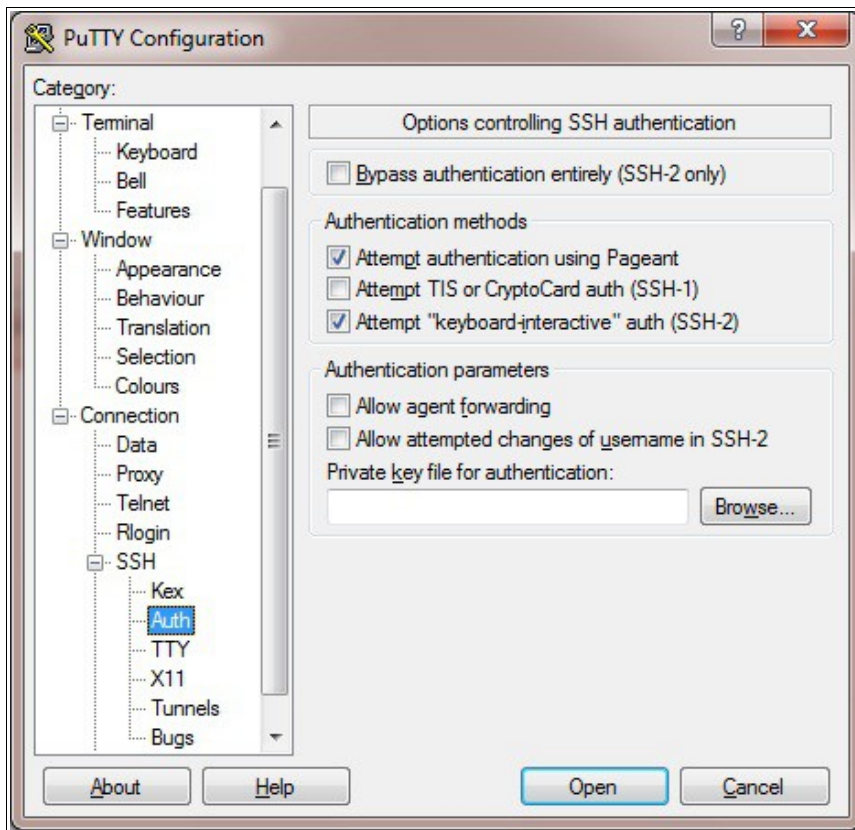
You can also choose to enter an additional pass-phrase for improved security, if you don't, the following message will be displayed:



Click **Yes** and save the file with the default .ppk extension

Now close PuTTYgen and start PuTTY

Expand the SSH section as shown below:



Click **Browse** and select the new .ppk file just created

When you open the SSH session, login as 'azureuser' – a password will no longer be required

Accessing the Appliance using SCP

Using Linux

First change the permission of the private key file to allow only the owner read access

```
# chmod 400 /path-where-saved/MyPrivateKey.key
```

Now start SCP specifying the private key file – login as user '**azureuser**'

```
# scp -i /path-where-saved/MyPrivateKey.key <local-file> azureuser@<LB-IP-address>:<remote-file>
```

or


```
# scp -i /path-where-saved/MyPrivateKey.key <local-file> azureuser@<LB-dns-name>:<remote-file>
```

N.B. to configure an FQDN in Azure under the Resource Manager model please refer to [this link](#)

Using Windows

Download WinSCP from: <http://winscp.net/eng/download.php>

With WinSCP, enter the relevant IP address, set the username field to **azureuser**, and browse to the private key file created previously using PuTTYgen as shown below:



Click **Login**

Configuration Examples

The following sections provide a number of examples to help illustrate how the load balancer can be deployed.

1) Load Balancing Web Servers - Single Subnet, Layer 7

This is a simple layer 7 example using one public subnet for both the load balancer and the web servers.

a) Setting up Azure

- Deploy the load balancer instance as described earlier
- Deploy your required web server VM's into the same VNet & subnet as the load balancer
- Ensure the Network Security Group includes port 80 (see page 8)
- Public IP addresses are not needed when deploying the web servers (real servers) instances since the load balancer is configured to send traffic to the private IP address of each web server

b) Setting up the Virtual Service

- Using the WUI, go to *Cluster Configuration > Layer 7 – Virtual Service* and click **[Add a New Virtual Service]**
- Enter the following details:

Label	Web-Cluster1	?	
Virtual Service	IP Address	10.0.0.22	?
	Ports	80	?
Layer 7 Protocol	HTTP Mode	▼	?
Manual Configuration	<input type="checkbox"/>		?

Cancel Update

- Enter an appropriate label for the VIP, e.g. **Web-Cluster1**
- Set the *Virtual Service IP address* field to the required IP address, e.g. **10.0.0.22**





*N.B. the current version of the appliance allows any IP address within the same subnet to be specified for the VIP, although the VIP will still only be accessible on the Private IP address allocated to the appliance. The public and private IP addresses can easily be viewed in the Azure Portal under **Network Interfaces** as shown below:*

NAME	PUBLIC IP ADDRESS	PRIVATE IP ADDRESS	SECURITY GR...
lbtest1135	40.113.125.93	10.0.0.4	lbtest1 ...

- Set the *Virtual Service Ports* field to **80**
- Leave *Layer 7 Protocol* set to **HTTP mode**
- Click **Update**

c) Setting up the Real Servers

- Using the WUI, go to *Cluster Configuration > Layer 7 – Real Servers* and click **[Add a new Real Server]** next to the newly created VIP
- Enter the following details:

Label	<input type="text" value="Web1"/>	
Real Server IP Address	<input type="text" value="10.0.0.23"/>	
Real Server Port	<input type="text" value="80"/>	
Weight	<input type="text" value="100"/>	

- Enter an appropriate label for the RIP, e.g. **Web1**
- Change the *Real Server IP Address* field to the required IP address, e.g. **10.0.0.23**
- Click **Update**
- Repeat the above steps to add your other Web Server(s)

d) Applying the new Layer 7 Settings

- Once the configuration is complete, use the **Reload HAProxy** button at the top of the screen to commit the changes

e) Testing & Verification

To test the configuration is working, browse to the public IP address or FQDN on port 80, i.e.

http://<Public IP>

or

http://<FQDN>

N.B. to configure an FQDN in Azure under the Resource Manager model please refer to [this link](#)

2) Load Balancing Web Servers - Single Subnet Layer 7 with SSL Termination

This is similar to the first example with the addition of setting up SSL termination on the load balancer. We generally recommend that SSL should be termination on the backend servers rather than the load balancer for scalability reasons, although in some cases terminating on the load balancer may be preferred.

a) Setting up Azure

- Deploy the load balancer instance as described earlier
- Deploy your required web server VM's into the same VNet & subnet as the load balancer
- Ensure the Network Security Group includes port 443 (see page 8)
- Public IP addresses are not needed when deploying the web servers (real servers) instances since the load balancer is configured to send traffic to the private IP address of each web server

b) Setting up the Virtual Service

- Using the WUI, go to *Cluster Configuration > Layer 7 – Virtual Service* and click **[Add a New Virtual Service]**
- Enter the following details:

Label	<input type="text" value="Web-Cluster1"/>	?	
Virtual Service	IP Address	<input type="text" value="10.0.0.22"/>	?
	Ports	<input type="text" value="80"/>	?
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	▼	?
Manual Configuration	<input type="checkbox"/>		?
<input type="button" value="Cancel"/> <input type="button" value="Update"/>			

- Enter an appropriate label for the VIP, e.g. **Web-Cluster1**
- Set the *Virtual Service IP address* field to the required IP address, e.g. **10.0.0.22**





*N.B. the current version of the appliance allows any IP address within the same subnet to be specified for the VIP, although the VIP will still only be accessible on the Private IP address allocated to the appliance. The public and private IP addresses can easily be viewed in the Azure Portal under **Network Interfaces** as shown below:*

NAME	PUBLIC IP ADDRESS	PRIVATE IP ADDRESS	SECURITY GR...
lbtest1135	40.113.125.93	10.0.0.4	lbtest1 ...

- Set the *Virtual Service Ports* field to **80**
- Leave *Layer 7 Protocol* set to **HTTP mode**
- Click **Update**

c) Setting up the Real Servers

- Using the WUI, go to *Cluster Configuration > Layer 7 – Real Servers* and click **[Add a new Real Server]** next to the newly created VIP
- Enter the following details:

Label	<input type="text" value="Web1"/>	
Real Server IP Address	<input type="text" value="10.0.0.23"/>	
Real Server Port	<input type="text" value="80"/>	
Weight	<input type="text" value="100"/>	

- Enter an appropriate label for the RIP, e.g. **Web1**
- Change the *Real Server IP Address* field to the required IP address, e.g. **10.0.0.23**
- Click **Update**
- Repeat the above steps to add your other Web Server(s)

d) Configuring SSL Termination

- Using the WUI, go to *Cluster Configuration > SSL Termination* and click **[Add a New Virtual Service]**
- Enter the following details:

Label	<input type="text" value="SSL-WEB"/>	?
Virtual Service IP address	<input type="text" value="10.0.0.22"/>	?
Virtual Service Port	<input type="text" value="443"/>	?
Backend Virtual Service IP Address	<input type="text" value="10.0.0.22"/>	?
Backend Virtual Service Port	<input type="text" value="80"/>	?
Ciphers to use	<input type="text" value="ECDH+AESGCM:DH+AES"/>	?
Do not insert empty fragments	<input checked="" type="checkbox"/>	?
SSL Terminator	<input type="radio"/> Pound <input checked="" type="radio"/> STunnel	?
Delay DNS Lookups	<input checked="" type="checkbox"/>	?
Disable SSLv2 Ciphers	<input checked="" type="checkbox"/>	?
Disable SSLv3 Ciphers	<input checked="" type="checkbox"/>	?
Allow Client Renegotiation	<input checked="" type="checkbox"/>	?
Disable SSL Renegotiation	<input checked="" type="checkbox"/>	?
Time To Close	<input type="text" value="0"/>	?
Set as Transparent Proxy	<input type="checkbox"/>	?

- Enter an appropriate label for the VIP, e.g. **SSL-WEB**
- Set the *Virtual Service IP address* to be the same as the VIP created in step c) e.g. **10.0.0.22**
- Set the *Virtual Service Ports* field to **443**
- Set the *Backend Virtual Service IP address* to be the same as the VIP created in step c) e.g. **10.0.0.22**
- Set the *Backend Virtual Service Ports* field to **80**
- Leave all other settings at their default values
- Click **Update**

SSL Certificate Notes:

- A default self-signed certificate will be used when setting up SSL Termination
- To change this, using the WUI, select: *Cluster Configuration > SSL Termination*
- Click [**Certificate**] next to the Virtual Service
- If you already have a certificate, use the **Upload prepared PEM/PFX file** option at the bottom of the screen to upload it
- If you don't have a certificate, you can create a CSR using the **Generate SSL Certificate Request** section. This will create the CSR in the upper pane of the **Upload Signed Certificate** section based on the settings you enter. This should be copied and sent to your CA

- Once the signed certificate is received copy/paste it (along with any required intermediate certificates) the lower pane of the **Upload Signed Certificate** section, and click **Upload Signed Certificate**

e) Applying the new Settings

- Once the configuration is complete, use the **Reload HAProxy** button at the top of the screen to commit the changes
- Once the configuration is complete, use the **Restart Stunnel** button at the top of the screen to commit the changes

f) Testing & Verification

To test the configuration is working, browse to the public IP address or FQDN on HTTPS port 443, i.e.

https://<Public IP>

or

https://<FQDN>

N.B. to configure an FQDN in Azure under the Resource Manager model please refer to [this link](#)

Testing & Validation

Testing Load Balanced Services

For example, to test a web server based configuration, add a page to each web servers root directory e.g. **test.html** and put the server name on this page for easy identification during the tests.

Use two or more clients to do the testing. Open up a web browser on each test clients and enter the URL for the VIP e.g. **http://104.40.133.119**

Provided that persistence is disabled, each client should see a different server name because of the load balancing algorithm in use , i.e. they are being load balanced across the cluster.

Why test using two clients? If you use a single client it will most likely keep on hitting the same server for multiple requests. This is to do with the way that the load balancing algorithms are optimized.

Diagnosing VIP Connection Problems

1. **Make sure that the device is active** - this can be checked in the WUI. For a single appliance, the status bar should report **Master & Active** as shown below:

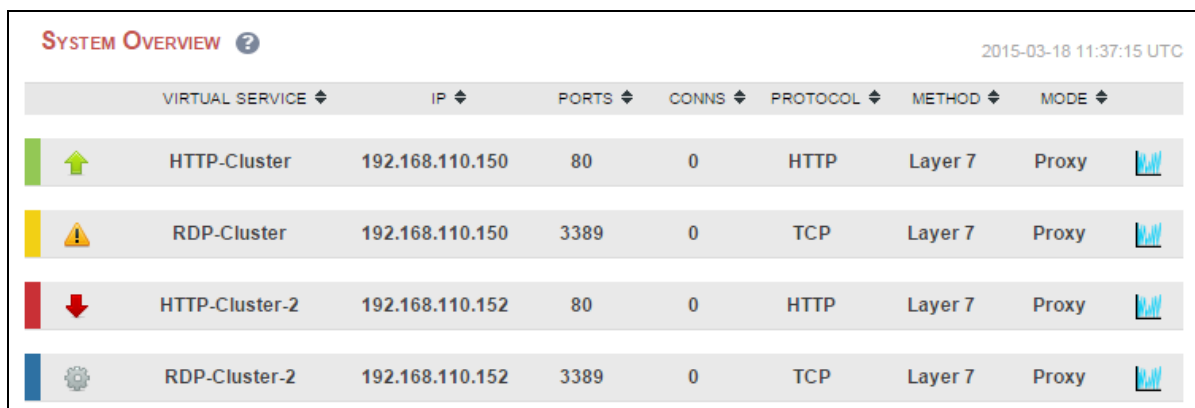


2. **Check that the VIP/floating IP is up** - Using *View Configuration > Network Configuration* verify that the VIP is active on the load balancer, if not check *Logs > Heartbeat* for errors.

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9001 qdisc pfifo_fast state UP qlen 1000
link/ether 02:bd:88:12:2f:5b brd ff:ff:ff:ff:ff:ff
inet 10.0.0.220/24 brd 10.0.0.255 scope global eth0
    valid_lft forever preferred_lft forever
inet 10.0.0.22/24 brd 10.0.0.255 scope global secondary eth0
    valid_lft forever preferred_lft forever
inet6 fe80::bd:88ff:fe12:2f5b/64 scope link
    valid_lft forever preferred_lft forever
```

The above example shows that the interface (10.0.0.220) and VIP address (10.0.0.22) are both up.

3. **Check that the Real Servers are up** - Using *System Overview* make sure that none of your VIPs are colored red. If they are, the entire cluster is down (i.e. all Real Servers). Green indicates a healthy cluster, yellow indicates that your cluster may need attention (one or more of the Real Servers may be down), and blue indicates all Real Server have been deliberately taken offline (by using either Halt or Drain).



The image shows a screenshot of the 'SYSTEM OVERVIEW' page. The table has columns for VIRTUAL SERVICE, IP, PORTS, CONNS, PROTOCOL, METHOD, and MODE. The status of each service is indicated by a colored bar on the left: green for healthy, yellow with a warning icon for attention, red with a down arrow for down, and blue with a gear icon for offline.

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy
RDP-Cluster	192.168.110.150	3389	0	TCP	Layer 7	Proxy
HTTP-Cluster-2	192.168.110.152	80	0	HTTP	Layer 7	Proxy
RDP-Cluster-2	192.168.110.152	3389	0	TCP	Layer 7	Proxy

4. Check the connection state -

Check *Reports > Layer 7 Status*. The default credentials required are:

username: loadbalancer
password: loadbalancer

This will open a second tab in the browser and display a statistics/status report as shown in the example below:

Statistics Report for pid 3261

> General process information

pid = 3261 (process #1, nbproc = 1)
 uptime = 0d 0h00m42s
 system limits: memmax = unlimited; ulimit-n = 81000
 maxsock = 80024; maxconn = 40000; maxpipes = 0
 current conns = 1; current pipes = 0/0; conn rate = 2/sec
 Running tasks: 1/5; idle = 100 %

Legend:
 active UP (green), active UP, going down (yellow), active DOWN, going up (orange), active or backup DOWN (red), backup UP (blue), backup UP, going down (purple), backup DOWN, going up (brown), not checked (grey), active or backup DOWN for maintenance (MAINT) (dark red).
 Note: UP with load-balancing disabled is reported as "NOLB".

Display option: Hide 'DOWN' servers, Refresh now, CSV export
 External resources: Primary site, Updates (v1.5), Online manual

Queue		Session rate			Sessions			Bytes		Denied		Errors		Warnings		Server											
Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle
Frontend	0	15	-	0	4	40 000	56	21 696	3 385 782	0	0	0	0	0	0	0	0	OPEN									
backup	0	0	-	0	0	-	0	0	0	0	0	0	0	0	0	0	0			1	-	Y					
RIP1	0	0	-	0	16	0	2	56	56	21 696	3 385 782	0	0	0	0	0	0	42s UP	L4OK in 0ms	1	Y	-	0	0	0s	-	
Backend	0	0	-	0	16	0	2	4 000	56	21 696	3 385 782	0	0	0	0	0	0	42s UP		1	1	1		0	0s	-	

Queue		Session rate			Sessions			Bytes		Denied		Errors		Warnings		Server											
Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle
Frontend	2	4	-	1	1	2 000	8	1 464	33 111	0	0	4	0	0	0	0	0	OPEN									
Backend	0	0	-	0	0	200	0	1 464	33 111	0	0	0	0	0	0	0	0	42s UP		0	0	0			0		

Taking Real Servers Offline

- Using the *System Overview* check that when you Halt one of the Real Servers the connections are redirected to the other server in the cluster.
- Stop the web service/process on one of the servers, wait a few seconds (for the load balancer to detect the change) and then refresh the browsers on both clients. They should now both switch to the same server (since one has been removed from the load balancing list). Also check that the server is shown red (down) in the system overview.
- Start the web service/process on the server, wait a few seconds and then refresh the browsers again. After a few refreshes they should again show different web servers. Also check that the server is shown green (up) in the system overview.

The *System Overview* shows the status as these tests are performed:

SYSTEM OVERVIEW 2015-04-30 08:35:41 UTC

VIRTUAL SERVICE	IP	PORTS	CONN	PROTOCOL	METHOD	MODE
HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy
REAL SERVER	IP	PORTS	WEIGHT	CONN		
RIP1	192.168.110.240	80	100	0	Drain	Halt
RIP2	192.168.110.241	80	0	0	Online (halt)	
RIP3	192.168.110.242	80	100	0	Drain	Halt

In this example:

'rip1' is green, this indicates that it's operating normally

'*rip2*' is blue, this indicates that it has been either Halted or Drained. In this example Halt has been used as indicated by *Online (Halt)* being displayed. If it had been drained it would show as *Online (Drain)*

'*rip3*' is red, this indicates that it has failed a health check

Using Reports & Log Files

The appliance includes several logs and reports that are very useful when diagnosing issues. Both are available as main menu options in the WUI. Details of both can be found in the administration manual.

Loadbalancer.org Technical Support

If you have any questions regarding the appliance don't hesitate to contact the support team support@loadbalancer.org or your local reseller.

For more details please refer to the administration manual:

<http://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf>

Appendix

1. Updating the Agent when Using newer Instance Sizes

- Connect to the instance using SSH (under Windows use PuTTY)
- If authentication was configured to use a username/password, login using these credentials
- If authentication was configured to use SSH keys, refer to page 14
- Once logged in, run the following command to gain root access:

```
$ sudo su
```

- now run the following commands to update and restart the agent:

```
# wget https://github.com/Azure/WALinuxAgent/archive/2.0.zip
```

```
# unzip 2.0.zip
```

```
# cd WALinuxAgent-2.0/
```

```
# cp /usr/sbin/waagent /usr/sbin/waagent.bak
```

```
# cp -f waagent /usr/sbin/waagent
```

```
# chmod 755 /usr/sbin/waagent
```

```
# service waagent restart
```

2. Company Contact Information

Website	URL : www.loadbalancer.org
North America (US)	Loadbalancer.org, Inc. 4250 Lancaster Pike, Suite 120 Wilmington DE 19805 USA Tel : +1 888.867.9504 Fax : +1 302.213.0122 Email (sales) : sales@loadbalancer.org Email (support) : support@loadbalancer.org
North America (Canada)	Loadbalancer.org Ltd 300-422 Richards Street Vancouver, BC V6B 2Z4 Canada Tel : +1 866.998.0508 Fax : +1 302.213.0122 Email (sales) : sales@loadbalancer.org Email (support) : support@loadbalancer.org
Europe (UK)	Loadbalancer.org Ltd. Compass House North Harbour Business Park Portsmouth, PO6 4PS UK Tel : +44 (0)330 3801064 Fax : +44 (0)870 4327672 Email (sales) : sales@loadbalancer.org Email (support) : support@loadbalancer.org
Europe (Germany)	Loadbalancer.org GmbH Alt Pempelfort 2 40211 Düsseldorf Germany Tel : +49 (0)211 9793 7203 Fax : +49 (0)30 920 383 6495 Email (sales) : vertrieb@loadbalancer.org Email (support) : support@loadbalancer.org